EONORYX

READ THE HORIZON

FOUNDER'S RESEARCH CHARTER & DISCOVERY FRAMEWORK 001

The Unnamed Transition

Emergence, persistence and transformation at the edge of what science can name

A foundational research charter establishing EONORYX as an independent, interdisciplinary frontier research initiative: its mission and method, a synthesis of four commissioned investigations, the Discovery Registry of candidate phenomena, eight priority research programmes, and a staged horizon to 2036

Steven Alber
Founder, EONORYX

November 2026
Version 1.0.1 · Foundational edition

IMPRINT

The Unnamed Transition

Founder's Research Charter & Discovery Framework 001. The foundational research publication of EONORYX, an independent interdisciplinary frontier research initiative investigating phenomena that humanity has not yet learned to name.

Commissioned and directed by	Steven Alber, Founder, EONORYX
Institution	EONORYX · eonoryx.com · "Read the Horizon"
Document class	Foundational research charter; version-controlled, revisable
Version	1.0.1 · November 2026 · first foundational edition, corrective release (Appendix E)
Companion files	Discovery Registry (machine-readable), Scientific Source & Verification Register, Brand & Publication Style Guide, Digital Research Platform Blueprint, editable master source
Preparation	An independently commissioned, AI-assisted scientific synthesis. The research methodology and assistance disclosure is given in Chapter 13 and Appendix D
Licence of cited works	All cited scientific works remain the property of their authors and publishers and are referenced for scholarly purposes

How to read this charter

This is a research charter, not a forecast and not a collection of findings. It states what EONORYX exists to investigate and how (Parts I and V); synthesizes four independent investigations into one architecture without pretending they agreed where they did not (Part II); registers the candidate phenomena that survived, each with a classification, evidence grade, competing explanations, falsifiable prediction and experiment (Part III); and sets out what will be done, in what order, with what resources, and what would count as failure (Part IV). Three rules govern every page: a name is a hypothesis, never a discovery; an evidence grade is attached to every claim and may go down as well as up; a negative result is published with the same prominence as a positive one.

CLASSIFICATION OF CANDIDATE PHENOMENA

CLASS I Emerging reality — supported by credible, peer-reviewed observations, not yet understood or unified.

CLASS II Theoretical frontier — grounded in known science, not yet experimentally demonstrated.

CLASS III Radical conjecture — would require revising existing assumptions.

CLASS IV Conceptual invention — a framework or category proposed by EONORYX investigations.

Classes are not rankings. An entry with two classes is a Class I observation with a Class IV framework proposed to explain it. Replication is recorded by the evidence grade, not the class; a Class I entry resting on a preprint alone is marked provisional.

EVIDENCE GRADES

E3 Replicated or independently confirmed.

E2 Bounded observation — peer-reviewed, or a preprint so labelled — not yet independently replicated or generalized.

E1 Precursors only — related findings exist, the specific claim is untested.

E0 No evidence; a proposal.

RESEARCH STATUS

PRIORITY enters a programme (Chapter 8). MONITOR tracked by the laboratory. HOLD awaits a partner or a proponent's specification. ~~ARCHIVED~~ rejected or superseded; retained in Chapter 7's archive.

Contents

I **Foundations**

01 The Founder's Vision

02 Why the future may be conceptually unrecognizable

03 The Unnamed Transition

04 The scientific foundations of emergence, persistence and transformation

II **Synthesis**

05 Synthesis of the four independent investigations

06 The EONORYX Discovery Taxonomy

III **Registry**

07 The Master Register of candidate phenomena

08 Eight priority research programmes

09 Evidence, mathematical frameworks and experimental falsification

IV **Programme**

10 The EONORYX Frontier Discovery Laboratory

11 The 90-day research programme

12 The 2027–2036 research horizon

V **Institution**

13 Scientific integrity, provenance and publication standards

14 EONORYX Digital Research Platform architecture

15 Founder's closing thesis — what humanity has not yet learned to ask

A–E Appendices: references and verification register; methodological notes; definitions; provenance; version history

PART I

Foundations

Why EONORYX exists, what it refuses to assume, and the three questions around which its research is organized.

CHAPTER 01

The Founder's Vision

EONORYX was founded on a suspicion that is easy to state and hard to act on: the most consequential change of the coming decade will not be a better version of anything we already have a word for. It will be a phenomenon that our categories were not built to see. This charter is a systematic attempt to look for it in a way that could fail.

Why an independent initiative

Large research institutions are organized by discipline, funded by milestone, and evaluated by publication within a field. These are good arrangements for deepening what is known. They are poor arrangements for noticing that two fields are measuring the same thing under different names, that a widely cited result does not say what it is taken to say, or that an observation nobody owns is more important than the ones everybody owns. The gaps between disciplines are precisely where an independent initiative can contribute, provided it holds itself to standards of evidence at least as strict as the institutions it stands outside.

EONORYX is small by design. It has no laboratory of its own, no appointed faculty and no completed research to its name. What it has is a method — stated in this charter — for turning a vague sense that "something is changing" into named hypotheses, each with a definition precise enough to be wrong, an experiment capable of testing it, and a criterion for abandoning it. Its product is not predictions. Its product is a register of well-posed questions, kept honestly, and the experiments that answer them.

What EONORYX investigates

Our object of study is the class of phenomena that may emerge over the next two to ten years at the intersections of artificial and biological intelligence, physics, materials, neuroscience, synthetic biology, communication and economic organization — phenomena for which humanity does not yet possess concepts, instruments or vocabulary. We are not interested in whether machines will become faster, networks denser or materials stronger; those trajectories are visible and well served. We are interested in whether a new *kind* of thing is appearing: a new carrier of cognition, a new form of memory, a new mode of organization, a new state of matter that behaves as a device, a new relation between knowing and understanding.

Four independent investigations were commissioned to search for such phenomena from different starting points. They are synthesized in Part II. Their most important shared result was not a list of candidates but a question that all four, working separately, arrived at in different words: **when the physical carrier of a capacity changes — when competence moves from one agent to many, from a brain to a tissue, from a processor to a material, from an institution to a protocol — what remains, and how would we measure it?** We call this the carrier question, and it organizes everything that follows.

Three dimensions

The charter organizes its research around three foundational dimensions, which Chapter 4 grounds in existing science:

- **Emergence.** Under what conditions do new capabilities, properties or forms of organization arise that are not present in, and not predictable from, the components? The hard version of this question asks for an order parameter — a quantity whose change marks the moment a collection becomes a unit.
- **Persistence.** How can information, competence, memory, identity or function survive the replacement of the physical or computational carrier that currently holds it? The hard version asks what invariant predicts survival.
- **Transformation.** Under what conditions does a system cross a threshold beyond which existing scientific categories become inadequate? The hard version asks how we would know, from inside the old categories, that this had happened.

These three dimensions are not a philosophy to which phenomena must conform. They are questions. Several entries in the Discovery Registry belong to more than one; some belong to none cleanly, and we say so. The architecture is designed to be revised by what it finds.

Principles

EONORYX commits to the following, and this charter is the instrument by which it can be held to them.

- **A name is a hypothesis.** Proposing a new term for a phenomenon is a tool for making it testable. It is never evidence that the phenomenon exists. Where a proposed term merely renames an established concept, we identify the precedent and withdraw the term.
- **Evidence is graded, publicly, and may go down.** Every claim carries a classification and an evidence grade. Grades are re-assessed on a fixed cadence and whenever a relevant result appears. The 2026 re-analyses that narrowed one of our own strongest candidates are treated as a model case, not an embarrassment.
- **Falsification is specified before data.** Every registered phenomenon carries a prediction that could be wrong and a criterion for closing the inquiry. Research programmes are preregistered. Kill criteria are stated in advance and honoured.
- **Negative results are publications.** A programme that closes a hypothesis has succeeded. The history of forecasting contains very few documented null results; EONORYX intends to add to them.
- **Attribution is exact.** Original researchers are credited for their findings. EONORYX does not claim experiments it did not run, instruments it does not have, or partners it has not secured. The methodology and assistance disclosure in Chapter 13 states how this document was produced.
- **Speculation, model, observation and discovery are kept apart.** The four are different things and the charter labels each.

What this charter is not

It is not a technology roadmap, and it makes no claim about which companies or countries will lead. It is not a manifesto for any particular future; Chapter 15 argues that the framing adopted here may itself be wrong and says what would show it. It is not a claim of priority: the literature search that supports it cannot establish that any idea here is globally new, and several are explicitly built on named predecessors. And it is not finished. Version 1.0 is a foundation. The Discovery Registry is a living document, and the version history in Appendix E is where the charter's own evolution will be recorded.

THE FOUNDING COMMITMENT

EONORYX does not predict the next version of the world. It investigates the phenomena from which an entirely different world may emerge — and it builds the instruments that could show, with equal clarity, that no such phenomena are there.

Steven Alber

FOUNDER, EONORYX · NOVEMBER 2026

CHAPTER 02

Why the future may be conceptually unrecognizable

Forecasting fails in characteristic ways, and the failures are not random. They follow from assumptions built into the method: that the future extends present categories, that it arrives as artifacts, that intelligence belongs to agents, that transformation is additive, and that what matters can be counted. This chapter names those assumptions, documents how long phenomena run ahead of their names, and states the discipline EONORYX applies to anomalies so that the search for the unrecognizable does not become a licence to see what it wishes.

2.1 The inverted telephone

In 1965 an observer looking at a rotary telephone could imagine a smaller telephone, a portable one, one with a screen. The forecast would have been correct on every axis it measured and would have missed the only thing that mattered. Being reachable, remembering, finding the way, knowing the price, proving who you are — these were verbs performed by a person. The device that eventually arrived did not improve the telephone; it moved those verbs out of the person and into an environment that performs them continuously. The decisive change was never the object. It was the relocation of capacities from their historical subject to a new carrier, and the object was simply the first visible place the relocation settled.

That gives EONORYX its search heuristic: *look for verbs leaving their subjects.* Ask not what new device will appear, but which capacities are loosening from the carriers in which we assume they reside. The four investigations synthesized in Part II were asked this question from different directions, and each found candidates. Several of them concern the same loosening seen from different sides.

2.2 Five systematic blind spots

Category extrapolation. Forecasts move along axes that already exist — faster, smaller, cheaper, more accurate — because those are the axes on which data exist. Every roadmap examined for this charter, from the international framework for sixth-generation mobile systems[1] to national digital strategies, is written in the categories of the present. That is a necessary condition for writing a roadmap and a sufficient condition for missing a category change.

Artifact bias. Forecasting expects the future to be made of things, because things can be photographed. But the transformations with the largest consequences have usually been relations and protocols: a packet-switching standard, a payment rail, a way of indexing pages. The word "Internet" was in use among engineers in 1974[2], decades before the public had anything to point at.

Agent-centrism. Forecasts about intelligence are forecasts about agents: how capable, how many, how aligned. The possibility that intelligence is better described as a property of populations, of materials, or of the structured world that learners converge upon is not on the axis and therefore not in the forecast. One investigation put this as a question of the *unit of analysis*: we count people, models, devices and firms, and attribute capability to them, when the capability may depend on an arrangement whose participants continually change.

Additive bias. Forecasts predict arrivals. They are poorly equipped to predict disappearances — of a boundary, a category, a scarcity. A method that only counts what appears will register a transformation late, after the boundary is already gone.

Monoculture. A newer failure mode, documented in 2024: when a research community adopts the same tools and the same tools shape which questions are asked, the community can produce more agreement while collecting fewer independent observations — a monoculture of knowing and a monoculture of knowers that generates illusions of understanding[3]. This matters for EONORYX in a specific way. Its four investigations were conducted with similar assistance; Chapter 5 therefore treats their agreements with suspicion as well as interest, and Chapter 13 makes independence of sources a formal requirement.

2.3 The naming lag

Phenomena run ahead of their names by a decade or more, and this is structural. The word "smartphone" was coined by a handset maker in 1997 for a concept device; only in the following decade did it become the ordinary name of a category[4]. "Social media" has a first citation in 1994 and did not overtake its rivals in public usage until 2009[5]. "Cloud computing" appears in an internal corporate strategy paper in 1996, in a chief executive's remarks in 2006 and in a newspaper in 2007[6]. In each case the phenomenon was operational, and in most cases already transformative, before the language caught up.

Two formulations describe the underlying problem. Collingridge's dilemma holds that a technology's effects cannot be known until it is widely used, by which time it is too entrenched to change[7]. Kuhn's account of anomalies holds that observations are not perceived as such until a framework exists to admit them. Together they imply that the most consequential phenomena of any decade are, during that decade, both unnamed and unseen. The right response is not better extrapolation but the deliberate construction of candidate categories and the experiments that would test them. The naming lag is the window in which an institution that treats naming as a research act — a hypothesis with a falsifier attached — can do original work.

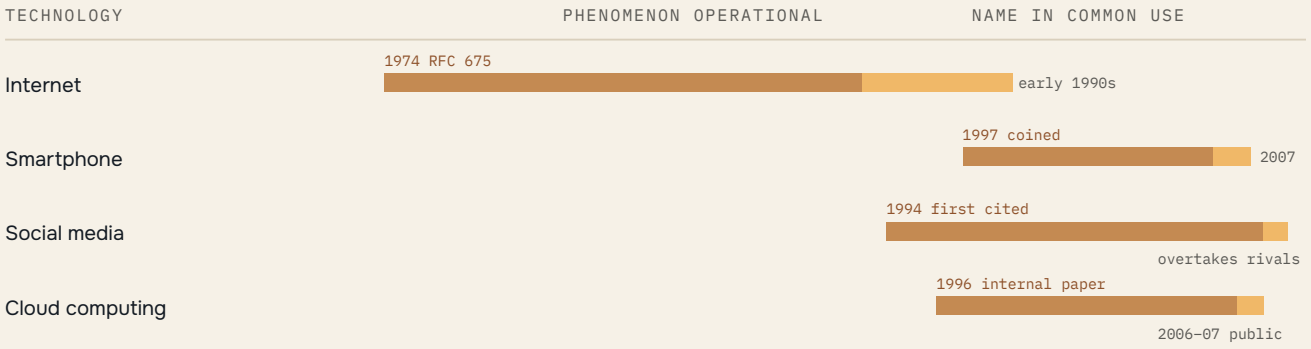


Figure 2.1 · The naming lag. Copper: the interval between first documented use of the term and its adoption as the common name of a category. Amber: the period of common use within the window shown. Sources: references cited in §2.3. The lag is ten years or more in every case.

2.4 The anomaly discipline

An institution that looks for the unrecognizable must guard most carefully against the reflex that treats every discrepancy as discovery. The last five years of physics supply the cautionary cases. The muon's anomalous magnetic moment was for two decades the most celebrated discrepancy between experiment and the Standard Model; by 2025 the measurement had reached its final precision and the theoretical prediction, recomputed with lattice methods, had moved to meet it, leaving a gap of about half a standard deviation and relocating the real disagreement inside theory[8]. A 2022 measurement of the W boson mass that differed from the Standard Model by seven standard deviations was contradicted by independent measurements consistent with the model[9]. Neither episode was a failure of science; both were failures of the reflex. Other anomalies persist without resolving — the disagreement between local and early-universe measurements of the cosmic expansion rate[10], and a preference in large-scale structure data for evolving dark energy whose significance moves with the data combination used[11][12] — and nothing in the next decade of technology turns on them.

RULE APPLIED THROUGHOUT THIS CHARTER

An observation enters the Discovery Registry as Class I only on credible, peer-reviewed publication, and is graded as replicated (E3) only once it has been independently replicated or confirmed; a single result, however striking, is E2 at most, and a result available only as a preprint is marked provisional. A mechanism is attributed only if a discriminating test exists that distinguishes it from the mundane alternatives — shared data, shared stimulus, estimator bias, measurement artifact, resource mismatch. A name is proposed only where an experiment could remove it. And where a proposed name renames an existing concept, the precedent is identified and the name withdrawn.

2.5 What would count as discovery

One of the four investigations proposed a distinction that this charter adopts as policy. A **new capability** is a performance improvement; it is the normal business of engineering and is not what EONORYX studies. A **new regime** is a reproducible change in how a system is organized or controlled — a threshold, a transition, a change in what predicts behaviour; most of the near-term opportunities in this charter are of this kind. A **new physical principle** requires residual evidence that survives the

strongest existing explanations; it cannot be manufactured by ambitious language, and no entry in this charter claims one. Keeping the three apart is the difference between a discovery institution and a producer of attractive futures.

CHAPTER 03

The Unnamed Transition

Four investigations searched independently for phenomena that humanity has not yet learned to name. None found a new law of nature, and this charter claims none. What they found, in four vocabularies, was the same movement: capacities that have always been located in specific carriers are becoming detachable from them. We call that movement the Unnamed Transition. This chapter states it as a hypothesis — with its evidence, its mechanisms and its strongest alternative — rather than as a conclusion.

3.1 The thesis

Four capacities have been bound, for the whole of human history, to four carriers. *Cognition* has resided in individual systems: a brain, later a program. *Knowledge* has resided in human witnesses: a result became knowledge when a person understood its justification. *Reality*, in the sense of a shared description of the world, has resided in a commons: the same newspaper, the same broadcast, the same street. And *computation* has resided in computers: a class of designed objects distinct from the matter around them. The Unnamed Transition is the hypothesis that all four bindings are loosening at once, and that the loosening is measurable in 2026.

Stated as the carrier question of Chapter 1: in each case a capacity is leaving the carrier that defined it, and the scientific task is to say what, if anything, remains invariant through the move — and how we would measure it.

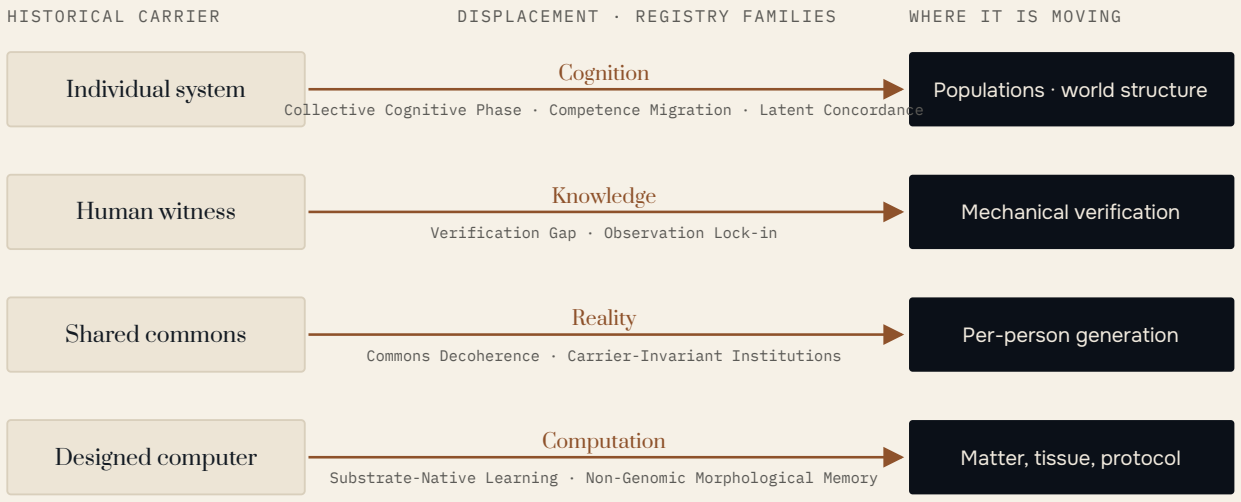


Figure 3.1 · The Unnamed Transition as four displacements. Each arrow is a hypothesis about a capacity leaving its historical carrier; the registry families named beneath each arrow (Chapter 6) are the testable forms of that hypothesis.

3.2 Displacement one — cognition leaving the individual system

Three observations from 2024–2026 do not fit the picture of intelligence as a property of a bounded system. Independently trained models — different architectures, modalities and data — converge toward a common relational structure of internal representation, and the convergence strengthens with capability[13]; embeddings from unrelated models can be translated into one another with no paired examples[14]; and re-analyses in 2026 indicate that what converges is local relational structure rather than global geometry[15][16], a narrowing that sharpens the claim rather than dissolving it. Populations of language-model agents left to interact form shared conventions and exhibit collective biases that no individual agent displays alone[17] — while other agent populations show no durable social structure at all[18], which tells us the phenomenon is conditional, not automatic. And a formal measure of information that exists only in the joint state of a group, long applied to brains[19], has now been applied to groups of artificial agents[20].

Three of the four investigations proposed, independently, an order parameter for the moment at which a set of interacting agents becomes a cognitive unit — synergistic information, the local learning coefficient of singular learning theory, and the probability of recovering a collective state after perturbation — and the fourth proposed proximity to criticality. Chapter 8 turns this convergence into a single programme that tests them against one another.

3.3 Displacement two — knowledge leaving the human witness

For the whole history of science a result was not knowledge until a person understood its justification. That condition is quietly ceasing to hold. In 2025 an evolutionary search system reported an algorithm for multiplying 4×4 complex matrices in 48 scalar multiplications, improving a bound that had stood since 1969; the result is verified by execution rather than by anyone's comprehension of why it works, independent mathematicians subsequently gave a 48-multiplication algorithm that avoids complex arithmetic[21], and the system's own authors report that it sometimes exploited weaknesses in the evaluator[22]. Machine-generated hypotheses in microbiology have been confirmed in the laboratory and published in peer-reviewed form[23][24]. In January 2026 a model solved an open problem in combinatorics and the proof was formalized within days; it was difficult enough to read that a human-readable translation had to be written[25]. A systematic attempt on hundreds of "open" problems found more than two-thirds of graded candidate solutions fundamentally flawed and warned of unintended plagiarism from the literature[26]. Verification has become mechanical; understanding has not become faster.

The quantity this creates has no name. We register it as the Verification Gap: the difference between the rate at which validated claims accumulate and the rate at which humans can reconstruct their justification. A related mechanism, proposed by another investigation, is adaptive observation lock-in: an automated scientific system that chooses what to measure so as to improve its own validation while losing the ability to detect what it has not modelled.

3.4 Displacement three — reality leaving the commons

A shared description of the world was always a by-product of shared media. Generative systems remove the by-product: every person can receive a unique artifact, tuned to them, in answer to the same question. The mechanisms are now measured. Deployed models affirm users' actions far more often than humans do and increase users' certainty and dependence[27]; real interaction memory

raises agreement with the user substantially[28]; the same question yields materially different facts for different personas across six models[29]; agent memory cuts response diversity by about a third[30]. What has not been measured is the divergence itself — what personalization does to the picture of the world people receive[31]. Two investigations arrived at the same reframing from different directions: the variable at stake is not polarization and not truth but *commensurability* — whether a population can still coordinate on a common description of what would happen if it acted together. One stated the testable form sharply: private predictive gains that coexist with a loss of joint control, at equal evidence and aligned objectives.

The same displacement has an institutional face. Coordination that once required a persistent intermediary — a platform holding the matching, the record and the remedy — is being specified in protocols that no single operator carries[32][33]. Whether the *meaning* of a permission, a commitment or a remedy survives the replacement of the operator that carried it is an open, testable question, and one investigation made it the centre of its programme.

3.5 Displacement four — computation leaving the computer

Beneath the three human-scale displacements the substrate itself is losing its boundary. Networks of self-adjusting resistors learn tasks through purely local physical updates without a processor[34], and elastic networks do the same mechanically[35]; the research group that built them describes the regime as physical networks that "become what they learn"[36]. Chips that compute by thermodynamic fluctuation rather than logic have been fabricated[37]. Cultured human neurons are sold as a computing substrate[38]. Biochemical networks inside single cells are modelled as capable of habituation[39]. Multicellular constructs assemble themselves from adult human cells and perform work no tissue was designed for[40]. And in 2026 a synthetic cell assembled from non-living components was reported, in preprint form, to replicate its genome, grow and divide[41]. "Computer" is becoming the name of a role that many kinds of matter can play rather than of a class of objects.

3.6 The strongest alternative

The thesis has a rival that the charter takes seriously and that Chapter 15 develops: that nothing is leaving anything, and that what looks like delocalization is *re-centralization through dependency*. Distributed agents may depend on a small set of models, identity services, energy systems and evaluation providers; interfaces may disappear while authority becomes harder to locate; apparently portable capabilities may fail the moment a shared provider is removed. On this reading the future is shaped not by capacities escaping their carriers but by carriers acquiring more intimate control over interpretation and action. The two readings make different predictions, and the research programmes in Chapter 8 are designed so that the dependency reading can win. A charter that could not lose this argument would not be worth writing.

STATUS OF THE THESIS

The Unnamed Transition is a Class IV conceptual framework resting on Class I observations. Its four displacements are hypotheses with distinct mechanisms and distinct evidence grades, recorded individually in the Discovery Registry. The charter does not assert that they are one phenomenon; it asserts that they share an experimental grammar — identify a function, specify its carrier, replace or disrupt the carrier, measure what survives, locate the information that explains survival — and that this grammar is worth building an institution around.

CHAPTER 04

The scientific foundations of emergence, persistence and transformation

The three dimensions of the EONORYX research architecture are not new ideas. Each has a scientific literature, a set of formal tools, a history of overclaiming, and a body of criticism that any serious programme must absorb before it begins. This chapter sets out, for each dimension, what is established, what is contested, which measures exist, and what EONORYX adds. It ends with the case for a fourth dimension that the evidence forced upon us.

4.1 Emergence

What is established

Emergence in its weak sense — properties of a whole that are not properties of its parts — is uncontroversial and everywhere: temperature, phase, traffic, markets. The scientifically interesting question is whether emergence can be made quantitative: whether there is a measurable sense in which a macroscopic description of a system carries information or causal power that the microscopic description does not, and whether that quantity changes sharply at identifiable thresholds. Three formal programmes address this.

Information decomposition. Partial information decomposition separates the information that a set of sources carries about a target into redundant, unique and synergistic components. Rosas and colleagues used it to define causal emergence as the presence of information about a system's future that is carried by a macroscopic variable and by no individual part, and showed that emergence in this sense requires the parts to be synergistic about their own future[42]; a review consolidated the approach[43]. Applied to the human brain, synergistic information concentrates in the prefrontal, parietal and default-mode regions that support high-level cognition and that expanded most in human evolution[19]. In 2025–26 the method was applied to groups of language-model agents: cross-agent synergy appears when agents are given distinct roles and prompted to model one another, and its link to group performance was marginal[20].

Causal emergence. Hoel and colleagues defined effective information as a measure of how much a system's causal structure determines its transitions, and proposed that a coarse-grained description can have higher effective information than the fine-grained one it summarizes[44]; a 2025 reformulation extends the framework[45]. The approach is contested: critics argue that the apparent causal gain of the macro-scale is an artifact of how the micro-scale is described and of unequal priors over interventions[46]. One of our investigations examined "causal-scale inversion" as a candidate phenomenon and rejected its novelty on exactly these grounds, retaining only the methods. This charter follows that judgement.

Singular learning theory. Watanabe's theory describes the geometry of statistical learning in models whose parameter-to-function map is degenerate — which includes neural networks — and shows that the free energy of such a model expands asymptotically with a leading correction governed by a

quantity called the real log canonical threshold or, in its local estimated form, the local learning coefficient[47]. Sharp changes in this coefficient during training have been proposed as signatures of phase transitions in what a network has learned, including the abrupt generalization known as grokking[48]. One investigation proposed the coefficient, computed over a *coupled* ensemble of agents, as the order parameter of a collective cognitive transition. It is a serious proposal with a serious cost: the estimator is expensive and its behaviour in multi-agent settings is untested.

Criticality. A long line of work holds that brains operate near a critical point between order and chaos, and that distance from criticality tracks cognitive state: conscious states have been associated with near-critical cortical dynamics[49], and anaesthetic loss of consciousness with increased distance from both avalanche criticality and the edge of chaos[50]. The hypothesis is attractive and the measures are mature — branching ratio, avalanche exponents, long-range temporal correlations, Lyapunov estimators — but the inference is treacherous: power laws arise without criticality[51], subsampling biases every estimator[52], and the field has warned itself repeatedly[53]. One investigation proposed criticality proximity as the order parameter of collective cognition; the charter retains it as a competitor to the other three.

What EONORYX adds

Four investigations independently proposed four candidate order parameters for the same transition. No one has run them against each other on the same system. The Emergence dimension of EONORYX is therefore organized around a single experimental object — an interacting population of learning agents whose interaction density can be swept — and a single question: which of these quantities, if any, changes discontinuously, predicts collective performance that individuals do not, and survives surrogate controls? This is Programme P1 in Chapter 8.

4.2 Persistence

What is established

Persistence of function through change of carrier has precedents in at least five fields, and the first task of any new work is to name them. In philosophy of mind, the extended-mind thesis holds that cognitive processes can be partly constituted by external resources[54], and distributed-cognition studies showed how competence resides in the arrangement of people and instruments rather than in any head[55]. In computer science, error correction is the engineering of information that survives the corruption of its carrier, and quantum error correction below threshold was demonstrated in 2024[56]. In developmental biology, Levin and colleagues have shown that a transient perturbation of bioelectric signalling in planarians can produce a lasting change in regenerative outcome — a two-headed animal whose fragments continue to regenerate two heads — a bounded, replicated observation[57] that the group interprets within a broader "bioelectric code" framework[58]. In organizational theory, institutions outlive their members by construction. And in physical learning, the question of whether a damaged adaptive network can recover a learned function without re-teaching is now askable because such networks exist[34].

What is contested

Each precedent has a weak and a strong reading, and the strong readings are contested. Extended cognition is accepted as a description and disputed as a metaphysics. The bioelectric code is accepted for bounded observations and disputed as a general mechanism; critics point to the difficulty of separating bioelectric state from metabolic state and to the gap between planarian and mammalian tissue. Institutional continuity is accepted as a fact and unexamined as a measurable invariant. None of these fields asks the question in the form one investigation posed it: *what must remain invariant for a capability to remain the same capability when every component that carries it can be replaced?*

What EONORYX adds

A shared experimental grammar. Identify a function defined independently of the system's self-description; specify a family of carrier replacements; measure the invariant that predicts continuity; find the falsifier where the invariant stops working. The grammar is the same whether the carrier is a set of agents (competence migration), a distributed network (recoverable collective state), a material (regenerative task memory), a tissue (morphological memory) or an operator (carrier-invariant institutions). Keeping the mechanisms distinct while sharing the grammar is the Persistence dimension's contribution, and Programmes P2, P5, P6 and P8 are its instances.

4.3 Transformation

What is established

Phase transitions in physics are the paradigm of a threshold beyond which the old description fails: the order parameter, the critical point, the divergence of susceptibility, universality. Non-equilibrium physics has extended the paradigm to systems that are driven and dissipative. Non-reciprocal interactions — where one species attracts a second that repels it — produce phases that have no equilibrium counterpart, organized around exceptional points of the dynamical matrix[59]. Time crystals break time-translation symmetry and have now been observed in semiconductors for tens of minutes[60], in liquid crystals as space-time crystals visible under a microscope[61], and as two-dimensional discrete time crystals on a 144-qubit processor[62]. In learning systems, grokking and related sharp transitions show that capability can change discontinuously with training. In social systems, convention formation in agent populations shows tipping points at which a committed minority flips a population's behaviour[17].

What is contested

Whether these are instances of one thing. "Phase transition" is used loosely across fields; the formal conditions for the term — a thermodynamic limit, a diverging correlation length, a universality class — are rarely met outside physics, and one investigation was careful to say that a finite network can show a useful operational transition without proving a singularity in an infinite-size limit. The charter adopts that care. Assembly theory, which proposes a measure of how much selection was required to produce an object[63], was offered by one investigation as a cross-substrate invariant of transformation; it has been sharply criticized on both its chemistry and its claims to novelty[64], and the charter records it as a contested tool rather than a foundation.

What EONORYX adds

Transformation, in the sense this charter needs, is not a property of a system but a relation between a system and the categories used to describe it: the moment at which the categories stop predicting. That gives a test. A transformation has occurred in a domain when a quantity that used to predict behaviour — an agent's individual score, a model's parameter count, a platform's take rate, a cell's genome — stops predicting it, and a different quantity, defined at a different scale or on a different carrier, starts. Each registry entry in this dimension names the old predictor and the candidate new one.

4.4 A fourth dimension — verification

The mandate for this charter allowed additional foundational dimensions if the evidence warranted them. One does. Across the four investigations, the most consistent concern was not about what systems can do but about what we can *know* about what they do: knowledge outrunning human understanding (the Verification Gap); automated science choosing observations that confirm its own models (observation lock-in); recursive training on synthetic data narrowing the variance of what is represented (model collapse, observed so far only in controlled settings[65], with roughly one in ten sampled web pages in mid-2026 showing signs of machine authorship[66]); and research monocultures producing agreement without evidence[3]. These are not emergence, persistence or transformation. They are about the health of the epistemic loop between a civilization and the world it is trying to describe — and they bear directly on EONORYX, whose own investigations were conducted with the systems in question.

We therefore adopt **Verification** as a fourth dimension: under what conditions does a knowledge-producing system — a science, an institution, a civilization — retain the ability to detect that it is wrong? Its registry family is the Epistemic Carrier Shift (Chapter 6); its programme is P4; and its institutional expression is Chapter 13, which applies the dimension to EONORYX itself.

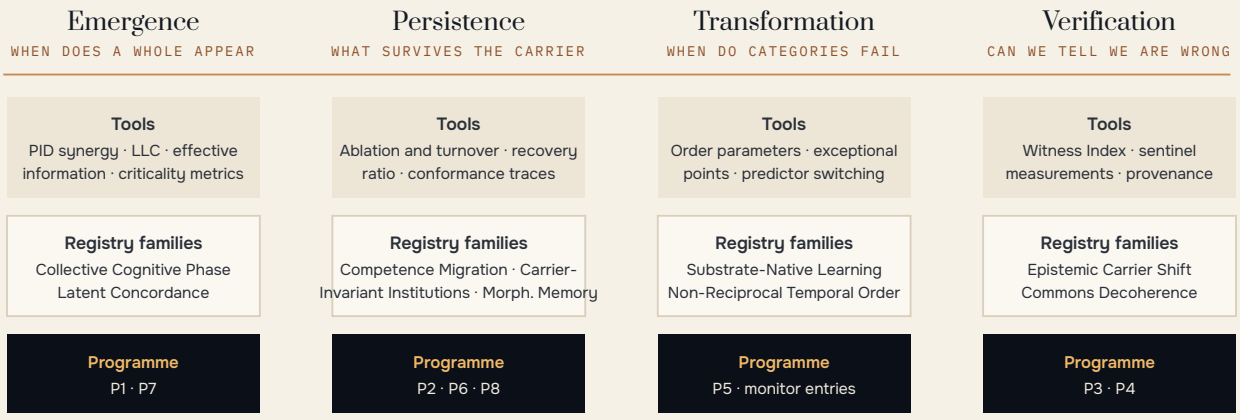


Figure 4.1 · The EONORYX research architecture: four dimensions, the formal tools each inherits, the registry families organized under each, and the priority programmes that test them. Several families belong to more than one dimension; the figure shows primary assignment.

4.5 What the architecture refuses to assume

It does not assume that the four dimensions are aspects of one law; Chapter 15 considers the possibility that they are four unrelated things and that the resemblance is an artifact of the method. It does not assume that any registered phenomenon exists. It does not assume that transformation is coming; the research programmes are built so that "nothing is happening" is a result they can return. And it does not assume that the categories it uses — order parameter, carrier, invariant, witness — will survive contact with the data. The architecture is a scaffold for experiments, and the experiments are allowed to take it down.

PART II

Synthesis

Four independent investigations, one research architecture. What they agreed on, where they contradicted one another, which terms rename old ideas, and what none of them supplied.

CHAPTER 05

Synthesis of the four independent investigations

Between October and November 2026 four research investigations were commissioned under an identical mandate: to search, without disciplinary boundaries, for phenomena that may emerge over the next two to ten years and for which humanity lacks concepts, instruments or vocabulary. Each was conducted with a different AI-assisted research system and received no sight of the others. This chapter treats them as four witnesses to the same question. It assesses each on evidence, originality, theoretical consistency, methodological quality and experimental feasibility; it then maps their agreements, contradictions and gaps. No investigation's conclusions are assumed correct, including the one that most closely resembles this charter's own framing.

5.1 The four witnesses

The investigations are referred to here as A, B, C and D. Their provenance — which system produced which, what the founder supplied, and how each was checked — is recorded in Appendix D and in the internal research archive. The public charter deliberately presents them as investigations rather than as a comparison of systems, because the object of interest is the phenomena, not the instruments.

INV.	ORGANIZING IDEA	STRONGEST ORIGINAL CONTRIBUTIONS	CHARACTER
A	Functional continuity under replacement: the unit of analysis is wrong	Competence migration; recovery probability as an order parameter in place of agreement; regenerative task memory in physical networks; counterfactual commons fracture (private gains, joint loss); institutional carrier invariance; adaptive observation lock-in; two executed negative controls	Methodologically the most disciplined; conservative on evidence; explicit about what it did not run
B	The substrate turn: matter, tissue and protocol negotiate order directly	Hysteretic matter intelligence (physical learning); bioelectric morphogenetic compilation; a singular-learning-theory order parameter for collective cognition; non-reciprocal active matter as self-clocked machinery; post-platform self-organization; three counter-theses (epistemic autophagy, hyper-synchrony, the apophatic turn)	The widest physical reach and the richest mechanisms; the most prone to overclaiming and to coined terminology; several citations needed correction
C	Dynamical regimes as cognitive variables	Criticality proximity as an order parameter; temporal experience multiplexing; bottom-up synthetic cell autonomy; neural intention alignment via error decoding; capacity-limited information flow	Represented only by an abbreviated reconstruction; its full text was not available. Entries are provisionally graded and flagged

INV.	ORGANIZING IDEA	STRONGEST ORIGINAL CONTRIBUTIONS	CHARACTER
D	Cognitive delocalization: capacities leaving their carriers	Latent concordance; machine normogenesis; a synergy-based order parameter; unwitnessed knowledge and the Verification Gap; consensus decoherence; closed-loop interneural integration staged with latency as the variable; infrastructure as instrument	Broadest source verification; strongest on the knowledge-and-reality displacements; weakest on physical substrates

5.2 Independent assessment

Evidence quality

Investigation A graded its own evidence most severely and was almost always right to. It rated two of its fifteen candidates at zero evidence and rejected one from the active portfolio; its Class I entries are bounded observations correctly described as such. Investigation D verified the largest number of sources against primary records and recorded disputes and 2026 re-analyses that narrowed its own strongest candidate. Investigation B cited the most frontier work but attached several claims to incorrect venues or titles — a 2026 ultrasound-interface paper that could not be located, a mechanical-network paper attributed to the wrong first author and an unverifiable journal locator, a DNA-computing paper given the wrong journal — and in two places treated a theoretical proposal as an established result. Investigation C's evidence cannot be assessed directly; where its summary names a specific project or number we checked it, with mixed results: the bottom-up synthetic cell it named is real and was reported in preprint in mid-2026[41], while the specific switching frequency it reported for rhythmic attention is not supported by the literature it points to, which places attentional sampling in the theta band[67][68].

Originality

Each investigation produced at least one idea that none of the others had. A's *recovery rather than agreement* as the order parameter of collective memory, and its *institutional carrier invariance*, are the sharpest reframings in the corpus. B's connection of *singular learning theory* to collective cognition, and its import of *non-reciprocal phase transitions* into morphogenesis, are the most ambitious interdisciplinary moves. C's *criticality proximity* gives the order-parameter question a fourth competitor from an established literature. D's *Verification Gap* is the only entry in the corpus that names a civilizational quantity and proposes to measure it. At the same time, every investigation coined at least one term for an existing concept; Chapter 6 audits them.

Theoretical consistency

The investigations are consistent in their rejections and divergent in their enthusiasms. All four rejected intelligence as a literal physical field. Three of the four (A, C and D) treated inter-brain coupling as unproven and in need of staged, controlled study; B treated inter-brain phase locking as "a reproducible laboratory reality," which overstates a literature whose own practitioners describe the evidence as correlational[69] and whose standard measures produce apparent coupling between brains that never interacted[70]. A and D converge on commons fragmentation as a coordination problem rather than a truth problem; B frames the same phenomenon as "ontological forking," which

is compatible but less testable. B and D converge on the inadequacy of platforms as long-term architecture; A asks the narrower and more answerable question of what survives when an operator is replaced.

Methodological quality

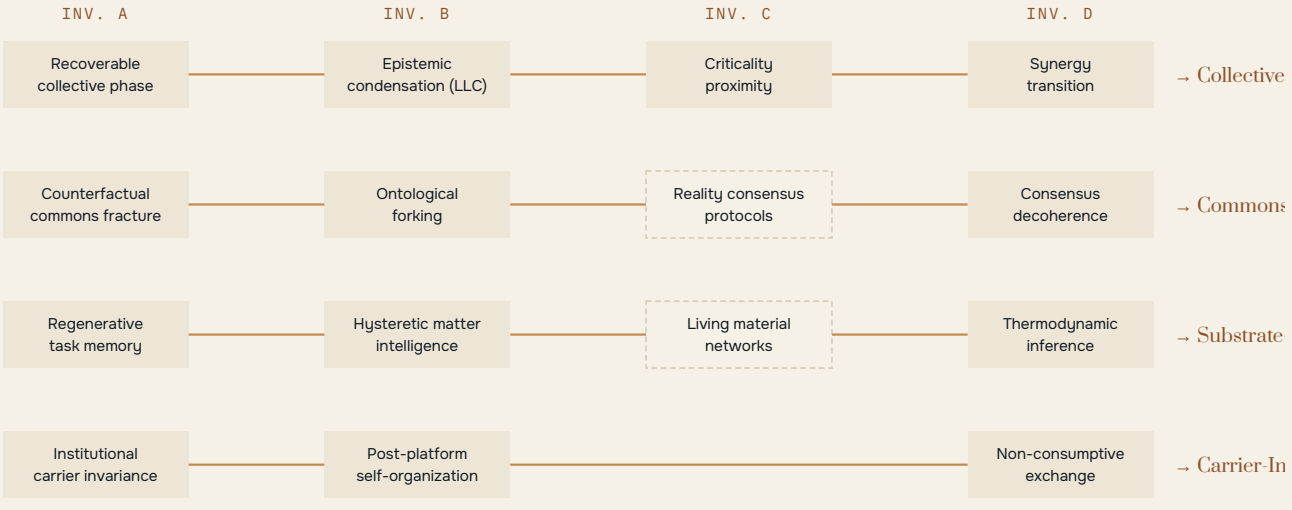
A supplied preregistration templates, a source register with evidence classes, and two executed, seeded, reproducible negative-control simulations that illustrate how collective-performance claims can be manufactured by independent components and by shared error — explicitly not tests of any hypothesis. D supplied a per-source verification log and discriminating-manipulation tables. B supplied numerical scores for novelty, plausibility, evidence and testability on a ten-point scale without stating how they were derived; the charter does not carry those scores forward. C's method cannot be assessed.

Experimental feasibility

For a small institution without a laboratory, feasibility is decisive. A's five programmes and D's four computational programmes can be run with public models, public data and rented compute. B's five programmes each end in a bench build with a university partner and a budget figure; those figures are stated as if secured and are not. C's one stated priority — validating its criticality measure on open electrophysiology and multi-agent time series against noise controls — is feasible and is folded into Programme P1.

5.3 Where the investigations agree

Independent agreement is the most valuable signal in the corpus, provided it is not an artifact of the shared mandate or of similar assistance. Three agreements survive that test because they arise in different vocabularies with different proposed measures.



Solid boxes: candidate stated with a definition and a test. Dashed: named in the abbreviated summary only. Blank: not proposed by that investigation. Each row became one registry family (Chapter 6); the four proposed measures in row 1 become competitors in Programme P1.

Figure 5.1 · Convergence map. Four rows of independent proposals that, on examination, concern the same phenomenon under different names and with different proposed measures.

Agreement one — an order parameter for collective cognition may exist and could be measured. A proposed the probability of recovering a collective state after perturbation; B proposed the local learning coefficient of singular learning theory computed over coupled agents, together with effective information; C proposed proximity to criticality; D proposed the synergistic fraction of task information. Four measures, one transition. None of the investigations noticed that it had three competitors. The charter's response is to make the competition the experiment.

Agreement two — the fragmentation of shared reality is a coordination phenomenon, not a truth phenomenon. A: private predictive gains with joint-control loss at equal evidence. D: declining mutual information between individuals' descriptions while individual accuracy is unchanged or improved. B: divergence into mutually untranslatable causal models. All three predict that the standard remedies — fact-checking, exposure diversity — aim at the wrong variable. A added a counter-weight the others lacked: experimental evidence that machine mediation can *increase* common ground[71], which makes fragmentation contingent rather than inevitable and turns the question into a comparative design.

Agreement three — computation is leaving the designed computer. B's physical learning, A's regenerative task memory, D's thermodynamic inference and living substrates, C's living material networks. The grades differ sharply: physical learning in resistor and elastic networks is replicated and peer-reviewed[34][36][72]; teacher-free recovery after damage is an untested extension; thermodynamic chips exist in prototype with efficiency claims unbenchmarked[37][73]; living-neuron substrates exist commercially with contested claims about what they do[74][38].

A fourth, weaker agreement — that the platform is not the long-term architecture of coordination — is real but divides on testability: A asks what survives operator replacement; B and D extrapolate to a post-platform regime. The charter keeps A's question as the testable core and the extrapolations as Class III.

5.4 Where they contradict one another

On inter-brain coupling. B treats dyadic phase locking as established and proposes closed-loop ultrasound neuromodulation to probe it; A grades the same phenomenon at zero evidence and D grades passive synchrony as disputed. The charter sides with A and D on the evidence and keeps B's proposed instrument as a possible phase-two manipulation, subject to the methodological critique that A and D both cite[70]. The paper B cited for the ultrasound interface could not be located; a 2026 review of ultrasonic brain–computer interfaces exists in preprint and is cited in its place[75].

On bioelectric memory. B asserts that anatomical form is "stored as a standing bioelectric voltage pattern" and that DNA "does not contain a geometric blueprint." A treats the same literature as a Class I bounded observation — persistent two-headed regeneration in planarians after a transient perturbation[57] — whose storage mechanism and generalization remain open, and notes the difficulty of separating bioelectric from metabolic state. D graded cellular learning as modelled but untested[39]. The charter adopts A's grading: the observation is real and replicated within one laboratory, the mechanism is the hypothesis, and independent replication of the persistent phenotype has not been located.

On causal emergence. B uses effective information as a load-bearing component of two hypotheses and of its single unifying question. A examined the same framework as a candidate phenomenon and rejected its novelty, citing the critique that macro-scale causal gains can be artifacts of how the

micro-scale is described[46]. The charter retains effective information as one *candidate measure* in Programme P1 and does not build on it.

On whether a transition is coming at all. A and D each end by arguing that the mandate's premise may be wrong and that "nothing is happening" must be a result the programme can return. B ends with a counter-thesis of its own — epistemic autophagy, the hyper-synchrony trap, and the possibility that the scarce category of 2036 is verified unmediated physicality — which is a different future rather than no transition. The charter adopts both: Chapter 15 treats "no transition" and "the opposite transition" as live alternatives with distinct predictions.

5.5 Redundancies and renamings

Two kinds of redundancy appear. Within the corpus, the same phenomenon carries up to four names (Figure 5.1); Chapter 6 resolves each into one registry family with the alternative names recorded as aliases. Against the literature, several coined terms rename an established concept; Chapter 6's precedent audit identifies the predecessor for each and withdraws the term where nothing is added. The most important cases: A's "causal-scale inversion" is causal emergence and A said so; B's "Hysteromimesis" names what its own cited authors already call physical learning; B's "Post-Platform Stigmergy" reuses a 1959 term from entomology for a process that B's own definition makes much more specific; D's "Latent Concordance" is a reading of the Platonic representation hypothesis rather than a new claim unless the data-disjoint experiment is run.

5.6 What none of them supplied

The gaps are as instructive as the findings. No investigation specified a *mechanism* by which synergy, learning-coefficient collapse, recovery and criticality would be related if they are measuring the same transition; P1 is designed to find out whether they are. No investigation supplied an independently replicated bioelectric-memory result outside a single laboratory. No investigation offered a theory of value for exchange among agents without consumption utility, though two named the problem. No investigation measured the divergence of information environments across real users; all measured ingredients. No investigation identified a verified case of a machine-produced result that humans have been unable to reconstruct; the Verification Gap is a prediction, not an observation. And no investigation supplied a benchmark on which its own order parameter could be compared with another's — which is why Chapter 8 begins by building one.

> THE SYNTHESIS IN ONE SENTENCE
>
> Four investigations, asked to find the unnamed, found the same four displacements under different names, proposed four rival measures for the first of them, disagreed most sharply where evidence was thinnest, and together left exactly the gaps that a small, disciplined institution is positioned to fill.

CHAPTER 06

The EONORYX Discovery Taxonomy

A taxonomy of hypotheses must do two things that pull against each other: group phenomena that are the same under different names, and keep apart phenomena that sound alike but differ in mechanism. This chapter defines the EONORYX taxonomy — nine families under four dimensions — and then audits every coined term in the corpus against the literature. Where a term renames an established concept, the precedent is named and the term withdrawn or demoted to an alias.

6.1 Principles of classification

- **Classify by mechanism and measure, not by vocabulary.** Two proposals belong to one family if they predict the same observable change by the same kind of mechanism, whatever they are called. They belong to different families if their proposed measures could disagree.
- **Every family has an experimental grammar.** A family is admitted only if a single experimental design could in principle test all its members. Families without such a design are not families; their members are listed individually.
- **Dimensions are primary assignments, not boxes.** A family is placed under the dimension whose question it most directly answers and cross-listed where it answers another.
- **Aliases are preserved.** Every name proposed by any investigation is recorded in the registry as an alias, so that provenance survives the merge.
- **Class and grade attach to entries, not families.** A family can contain a Class I observation and a Class III extrapolation; the taxonomy does not average them.

6.2 The nine families

#	FAMILY	DIMENSION	SHARED EXPERIMENTAL GRAMMAR	ENTRIES (CHAPTER 7)
F1	**Collective Cognitive Phase**	Emergence	Sweep the interaction density of a population of learning agents; measure candidate order parameters and collective performance; test for a discontinuity that individuals do not predict and surrogates do not reproduce	DR-01 to DR-05
F2	**Latent Concordance**	Emergence	Train independent learners on disjoint data from worlds of graded structural distance; measure relational alignment against calibrated nulls	DR-06
F3	**Competence Migration**	Persistence	Learn a task; replace participants and/or reset environment in a crossed design; locate by intervention where the task-relevant information resides	DR-07, DR-08

#	FAMILY	DIMENSION	SHARED EXPERIMENTAL GRAMMAR	ENTRIES (CHAPTER 7)
F4	**Carrier-Invariant Institutions**	Persistence · Transformation	Execute a bounded institutional process on independent implementations; replace the operator mid-process; measure semantic divergence and unauthorized actions	DR-09 to DR-11
F5	**Non-Genomic Morphological Memory**	Persistence	Perturb a non-neural tissue's physiological state transiently; measure persistence of altered developmental outcome with orthogonal readouts and blinded phenotyping	DR-12 to DR-14
F6	**Substrate-Native Learning**	Transformation	Implement learning in a physical, thermodynamic or living substrate; test for properties no digital implementation has — teacher-free recovery, energy crossover, interests	DR-15 to DR-19
F7	**Non-Reciprocal Temporal Order**	Transformation	Drive a system across a non-reciprocal exceptional point or into a time-crystalline phase; test whether temporal order is usable as a device property beyond the noise it adds	DR-20, DR-21
F8	**Epistemic Carrier Shift**	Verification	Measure the ratio of validated to witnessed claims; measure whether adaptive observation reduces detection of unmodelled causes; measure representational variance under recursive training	DR-22 to DR-25
F9	**Commons Decoherence**	Verification · Transformation	Randomize personalization of interpreters at equal evidence; measure individual accuracy, pairwise commensurability and joint control separately	DR-26, DR-27
—	**Unaffiliated entries**	various	Entries that passed the evidence rule but share no grammar with a family	DR-28 to DR-35
—	**Archive**	—	Rejected or superseded proposals, retained with the reason	DR-A1 to DR-A9

6.3 How the four investigations map onto the families

The table records, for each family, which investigation proposed a member and under what name. Names in the corpus are preserved as aliases in the registry.

FAMILY	INVESTIGATION A	INVESTIGATION B	INVESTIGATION C	INVESTIGATION D
F1 Collective Cognitive Phase	Recoverable collective phase	Noo-Phase Condensation / Symbiotic Epistemic Condensation	Criticality proximity; collective cognitive structural phase transitions	Synergy Transition; Machine Normogenesis
F2 Latent Concordance	—	Cross-Substrate Ontological Convergence	—	Latent Concordance

FAMILY	INVESTIGATION A	INVESTIGATION B	INVESTIGATION C	INVESTIGATION D
F3 Competence Migration	Competence migration; endogenous problem-space growth	—	—	(implicit in cognitive delocalization)
F4 Carrier-Invariant Institutions	Institutional carrier invariance	Autopoietic Protocol Metabolism / Post-Platform Stigmergy	—	Non-Consumptive Exchange; Model-Synchronous Communication
F5 Non-Genomic Morphological Memory	Developmental target persistence; assembly-released organismality	Electro-Morphostasis / Bioelectric Morphogenetic Compilation	Synthetic cell-cycle autonomy (adjacent)	Somatic Computation
F6 Substrate-Native Learning	Regenerative task memory	Hysteromimesis; Thermodynamic Fluctuation Inference; Chemoputational wetware	Living material transistor networks; substrate-agnostic computation boundaries	Thermodynamic Inference; Hybrid Cognitive Substrates
F7 Non-Reciprocal Temporal Order	—	Asymmetric Teleo-Matter; Boundary Time-Crystal Metrology	—	Chronomatter
F8 Epistemic Carrier Shift	Adaptive observation lock-in	Epistemic Autophagy (counter-thesis)	Capacity-limited information flow (adjacent)	Unwitnessed Knowledge / Verification Gap; Autogenic Software
F9 Commons Decoherence	Counterfactual commons fracture	Ontological Forking / Consensual Reality Perceptual Compilation	Personalized reality consensus protocols	Consensus Decoherence

6.4 The precedent audit

The mandate for this charter requires that where a proposed term merely renames an established scientific concept, the precedent be identified. The audit below applies a single test to every coined term in the corpus: *does the term pick out a measurable claim that the precedent does not?* Three outcomes are possible. **Withdrawn**: the term adds nothing and the precedent's name is used. **Alias**: the term is recorded for provenance but the registry entry carries a plainer name. **Retained**: the term names a claim with no adequate existing name, and the claim is stated.

COINED TERM (SOURCE)	PRECEDENT	WHAT, IF ANYTHING, THE TERM ADDS	OUTCOME
Hysteromimesis / Hysteretic Matter Intelligence (B)	"Physical learning" and "coupled learning" — the terms used by the authors of the cited experiments[34] [36]; earlier "physical neural networks"[76]	Nothing beyond the precedent; B's own definition is the Durian–Liu group's regime restated	Withdrawn → Physical learning networks
Regenerative task memory (A)	Robustness and fault tolerance in learning systems; self-healing materials	A specific, untested claim: recovery of a learned relation after damage *without teacher replay*, measured by a defined recovery ratio	Retained
Noo-Phase Condensation / Symbiotic Epistemic Condensation (B)	Phase transitions in learning (grokking); the local learning coefficient[47]; collective intelligence	A specific measure applied to a new object (coupled agents); the name adds nothing to the measure	Alias of DR-02 (LLC order parameter)
Synergy Transition (D)	Information-decomposition emergence[42]; synergistic core[19]	The prediction of a discontinuity in interaction density; the measure is the precedent's	Alias of DR-01 (synergy order parameter)
Recoverable collective phase (A)	Distributed error correction; attractor memory; norm emergence	Recovery probability as the order parameter in place of agreement; a stated finite-size protocol	Retained
Criticality proximity (C)	The critical brain hypothesis and its distance measures[49][50]	Nothing new as a measure; new as a competitor applied to agent populations	Alias of DR-04
Machine Normogenesis (D)	Norm emergence in multi-agent systems; convention formation[17]	Nothing beyond the precedent	Withdrawn → Convention formation in agent populations
Latent Concordance (D); Cross-Substrate Ontological Convergence (B)	The Platonic representation hypothesis[13] and its 2026 refinements[15]	A causal claim — that convergence is controlled by world structure rather than model lineage — testable only by the data-disjoint design; until run, the term is a reading of the precedent	Retained provisionally; falls to alias if the experiment is not run by end-2027
Competence migration (A)	Extended cognition[54]; distributed cognition[55]; external memory	An interventional map of where competence moves during training and turnover; the precedents name a location, this names a measured change in causal dependence	Retained
Institutional carrier invariance (A)	Interoperability; portable identity; protocol specification[77]	A specified invariant — the operational meaning of permission, commitment, revocation and remedy — tested across operator replacement	Retained

COINED TERM (SOURCE)	PRECEDENT	WHAT, IF ANYTHING, THE TERM ADDS	OUTCOME
Autopoietic Protocol Metabolism / Post-Platform Stigmergy (B)	Stigmergy (Grassé, 1959); autopoiesis (Maturana & Varela); decentralized exchange; state channels	An extrapolation that combines three precedents; no measure beyond A's; "metabolism" is metaphor	Alias of DR-10 (Class III extrapolation)
Non-Consumptive Exchange (D)	Mechanism design; algorithmic trading; agent-based computational economics	Names a genuine gap — a theory of value for parties without consumption utility — but supplies no theory	Retained as a question, not a phenomenon
Electro-Morphostasis / Bioelectric Morphogenetic Compilation (B)	Developmental bioelectricity; the "bioelectric code"[58]; morphostasis (an existing term in developmental biology)	Nothing beyond Levin's own framing; "compilation" is metaphor	Withdrawn → Bioelectric morphological memory (Levin's framing, graded as A graded it)
Developmental target persistence (A)	The planarian two-headed phenotype[57]	An operational definition that avoids claims about mechanism or intention	Retained
Asymmetric Teleo-Matter (B); Chronomatter (D)	Non-reciprocal phase transitions[59]; time crystals; active matter	"Teleo" imports purpose without evidence; "Chronomatter" adds an engineering framing (temporal order as a device property) but no measure	Both withdrawn → Non-reciprocal temporal order
Unwitnessed Knowledge / Verification Gap / Witness Index (D)	Philosophy of computer-assisted proof; epistemic dependence	A civilizational quantity with a rate and a proposed measurement protocol; no precedent names the stock or proposes to measure it	Retained
Adaptive observation lock-in (A)	Confirmation bias; optimal experimental design; scientific monocultures[3]	A specific mechanism — the missing evidence is never collected — with a synthetic-world test	Retained
Epistemic Autophagy (B)	Model collapse[65]	Nothing beyond the precedent, extended rhetorically to civilization	Withdrawn → Model collapse (monitor entry)
Counterfactual commons fracture (A); Consensus Decoherence (D); Ontological Forking (B)	Filter bubbles; polarization research; epistemic fragmentation	A and D add the dissociation claim (private accuracy up, joint control down) and a measure (joint regret; pairwise mutual information). B adds a label	Merged → Commons Decoherence (one entry, two measures); B's term an alias

COINED TERM (SOURCE)	PRECEDENT	WHAT, IF ANYTHING, THE TERM ADDS	OUTCOME
Causal-scale inversion (A)	Causal emergence[44]	Nothing; A itself rejected the novelty	Withdrawn; methods retained
Field-mediated cognition / Intelligence field (A, mandate)	None — no mechanism, carrier or instrument	Nothing testable	Archived
Intention Capture (D); Homeostatic Pre-Intentional Allostasis (B)	Anticipatory and proactive interfaces; readiness potentials; allostasis	A measurable criterion (decode-before-report latency) under D's name; B's term adds metaphor	Retained as Intention Capture; B's term an alias
Parasitic Measurement (D)	Opportunistic sensing; distributed acoustic sensing[78]	A general principle and a governance question; no new measure	Alias → Infrastructure as instrument

6.5 What the audit changes

The audit covers twenty-three proposals, several of which arrived under more than one coined name. Counted by proposal: seven are withdrawn or archived in favour of the established term, five are demoted to aliases of a registry entry, ten are retained because they name a measurable claim that the literature does not, and one group of three is merged into a single entry under a new name. Counted by name, about a third of the coined names are withdrawn outright and roughly another third survive only as aliases. The registry in Chapter 7 uses the retained names and the plain descriptive names; the withdrawn and alias names appear only in each entry's aliases field, so that provenance is preserved without the vocabulary. This is the practical meaning of the principle that a name is a hypothesis: the audit is the first experiment the names have been subjected to, and about a third of them did not survive it.

PART III

Registry

The Master Register of candidate phenomena, the eight priority programmes that test them, and the evidence, mathematics and falsifiers behind each.

CHAPTER 07

The Master Register of candidate phenomena

The EONORYX Discovery Registry, version 1.0.1. Thirty-five active entries in nine families and an unaffiliated group, and nine archived proposals retained with the reason for their rejection. Each active entry carries a unique identifier, a definition precise enough to be wrong, its provenance across the four investigations, the scientific precedents it builds on, a proposed mechanism, a classification and evidence grade, the competing explanations that must be excluded, falsifiable predictions, a proposed experiment, resource requirements, the earliest meaningful milestone, and a research status. Archived entries carry the same descriptive fields, and in place of an experiment, a budget and a milestone — none is assigned to an archived entry — they carry the condition under which they would be reopened. One Class I entry rests on a preprint alone and is marked provisional. The register is published in machine-readable form alongside this charter and is versioned; entries are re-graded on a fixed cadence. No entry is a confirmed discovery. Several are established findings of other researchers, included because they are the experimental substrate or the precedent of an EONORYX hypothesis, and they are labelled as such.

7.1 Summary

ID	NAME	CLASS	EV.	STATUS	PROGRAMME	EARLIEST MILESTONE
F1 · COLLECTIVE COGNITIVE PHASE · EMERGENCE						
EON-DR-01	**Synergy order parameter for collective cognition**	II / IV	E1	Priority	P1	Q2 2027: first order-parameter tournament result (positive or null).
EON-DR-02	**Learning-coefficient order parameter for coupled learners**	II	E1	Priority	P1	Q2 2027: coefficient trajectories on the tournament benchmark.
EON-DR-03	**Recovery order parameter for collective memory**	II	E1	Priority	P1	Q2 2027: finite-size recovery curves.
EON-DR-04	**Criticality-proximity order parameter**	II	E1	Priority	P1	Q2 2027: criticality measures on the tournament benchmark with controls.
EON-DR-05	**Convention formation in agent populations**	I	E2	Monitor	P1 (benchmark substrate)	Q1 2027: reproduction across two model families.
F2 · LATENT CONCORDANCE · EMERGENCE						

ID	NAME	CLASS	EV.	STATUS	PROGRAMME	EARLIEST MILESTONE
EON-DR-06	**Latent Concordance**	I / IV	E2	Priority	P7	Q3 2027: data-disjoint alignment result; the term lapses to an alias if not tested by end-2027.
F3 · COMPETENCE MIGRATION · PERSISTENCE						
EON-DR-07	**Competence migration**	IV	E1	Priority	P2	Day 30 of the 90-day programme: ablation matrix, even if every apparent migration is explained by ordinary memory.
EON-DR-08	**Endogenous problem-space growth**	II	E1	Monitor	—	2028: feasibility study after P2 results.
F4 · CARRIER-INVARIANT INSTITUTIONS · PERSISTENCE · TRANSFORMATION						
EON-DR-09	**Institutional carrier invariance**	IV	E1	Priority	P6	Day 45: conformance benchmark v0 and failure corpus.
EON-DR-10	**Post-platform self-organization**	III	E1	Monitor	P6 (boundary-condition simulation)	2027: boundary-condition map; promotion to Class II only if conditions are found and reproduced.
EON-DR-11	**Non-consumptive exchange**	II / IV	E1	Monitor	P6 (theory track)	2027: a formal statement of the valuation problem; data analysis contingent on access.
F5 · NON-GENOMIC MORPHOLOGICAL MEMORY · PERSISTENCE						
EON-DR-12	**Developmental target persistence**	I	E2	Priority	P8	Q1 2027: audit and simulation; 2028: partner replication decision.
EON-DR-13	**Assembly-released organismality**	I	E2	Monitor	—	Monitor for independent replication and for in-vivo results.
EON-DR-14	**Non-neural cellular learning**	II	E1	Monitor	P8 (literature)	2027: first wet-lab tests of the habituation models.
F6 · SUBSTRATE-NATIVE LEARNING · TRANSFORMATION						
EON-DR-15	**Physical learning networks**	I	E3	Monitor	P5 (foundation)	Re-graded on any failed replication or retraction; otherwise reviewed at the annual revision.
EON-DR-16	**Regenerative task memory**	II	E1	Priority	P5	Day 60: feasibility or impossibility result in simulation.
EON-DR-17	**Thermodynamic inference hardware**	II	E2	Monitor	—	2027: first independent benchmark.
EON-DR-18	**Living neural substrates**	I	E2	Monitor	—	2027: independently reproduced benchmark.
EON-DR-19	**Chemical and molecular reaction computing**	I	E2	Monitor	—	Promotion trigger: independent replication reported.
F7 · NON-RECIPROCAL TEMPORAL ORDER · TRANSFORMATION						

ID	NAME	CLASS	EV.	STATUS	PROGRAMME	EARLIEST MILESTONE
EON-DR-20	**Non-reciprocal temporal order**	I	E3	Monitor	—	2028: a demonstrated device property.
EON-DR-21	**Dissipative time-crystal sensing**	II	E1	Monitor	—	Monitor for first experiment.
F8 · EPISTEMIC CARRIER SHIFT · VERIFICATION						
EON-DR-22	**The Verification Gap**	I / IV	E1	Priority	P4	Day 90: first Witness Index for formal mathematics and algorithm design.
EON-DR-23	**Adaptive observation lock-in**	II	E1	Priority	P4 (companion)	Day 60: simulation result.
EON-DR-24	**Model collapse in deployed systems**	II	E2	Monitor	—	Promotion trigger: a measured loss of representational variance in a deployed model, traceable to recursive training data, reported in peer review.
EON-DR-25	**Autogenic software lineages**	II	E2	Monitor	—	2027–28.
F9 · COMMONS DECOHERENCE · VERIFICATION · TRANSFORMATION						
EON-DR-26	**Commons Decoherence**	I / IV	E1	Priority	P3	Day 45: audit result; day 75: human-panel interim read.
EON-DR-27	**Intention Capture**	II	E1	Monitor	—	2027: decode-before-report in a consumer-grade loop, if any.
F0 · UNAFFILIATED ENTRIES · VARIOUS						
EON-DR-28	**Closed-loop interneural integration**	III	E1	Hold	Laboratory protocol track	Day 90: validated protocol published regardless of outcome; partner decision.
EON-DR-29	**Infrastructure as instrument**	I	E3	Monitor	—	Review point: the predicted new-domain detection within the year; promotion to a programme only if the governance question becomes actionable.
EON-DR-30	**Bottom-up synthetic cell autonomy**	I (prov.)	E2	Monitor	—	2027: peer review and replication status.
EON-DR-31	**Coherent evidence accessibility**	I	E2	Monitor	—	Promotion trigger: advantage reproduced under realistic noise and readout by an independent group.
EON-DR-32	**Certified indefinite causal order**	I	E3	Monitor	—	Review point: a certification of indefinite causal order with the detection loophole closed.
EON-DR-33	**Rhythmic perceptual multiplexing**	I	E3	Monitor	—	Reviewed at the annual revision; no EONORYX milestone.

ID	NAME	CLASS	EV.	STATUS	PROGRAMME	EARLIEST MILESTONE
EON-DR-34	**Coherent processes in neural proteins**	III	E2	Monitor	—	2028: coherence times at functional timescales, if any.
EON-DR-35	**Neural intention and error alignment**	I	E2	Monitor	—	Promotion trigger: independent replication reported.
FA · ARCHIVE — REJECTED OR SUPERSEDED · NOT ASSIGNED (ARCHIVED)						
EON-DR-A1	**Intelligence as a physical field / field-mediated cognition**	III	E0	Archived	—	—
EON-DR-A2	**Passive intersubjective coupling**	III	E0	Archived	—	—
EON-DR-A3	**Causal-scale inversion**	IV	E2	Archived	—	—
EON-DR-A4	**Retrocausal quantum error protection**	III	E0	Archived	—	—
EON-DR-A5	**Capacity-limited information flow (Page-curve analogy)**	III	E0	Archived	—	—
EON-DR-A6	**Reality Negotiation as a product category**	IV	E0	Archived	—	—
EON-DR-A7	**Scale-free bioelectric–mycelial ecological percolation**	III	E0	Archived	—	—
EON-DR-A8	**'General intelligence' as a phenomenon**	IV	E0	Archived	—	—
EON-DR-A9	**Model-synchronous communication as a distinct phenomenon**	II	E1	Archived	—	—

7.2 Active entries

F1 Collective Cognitive Phase · Emergence

EON-DR-01 Synergy order parameter for collective cognition

CLASS II | CLASS IV | E1 · PRECURSORS ONLY | PRIORITY · P1

Definition. A population of interacting learning agents undergoes a collective cognitive transition at interaction density ρc if the synergistic component of information about a task variable, computed by partial information decomposition over the agents' states, changes discontinuously (or its susceptibility peaks) at ρc while the redundant component does not; above ρc, task-relevant information is carried only by the joint state.

PROVENANCE
Investigation D (hypothesis); measure from Rosas et al. 2020 and Luppi et al. 2022; first LLM-group application Riedl 2026.

ALIASES
Synergy Transition (D)

PRECEDENTS
Partial information decomposition (Williams & Beer 2010); causal emergence via synergy (Rosas et al. 2020); synergistic core of the human brain (Luppi et al. 2022).

PROPOSED MECHANISM
Coupling plus memory creates conditional dependence between agents' outputs and the joint history; below a critical coupling, task information is redundantly held by individuals; above it, symmetry breaking locks the population into states no individual determines and task information migrates into correlations.

NOVELTY ASSESSMENT
The measure is not new. The prediction of a discontinuity in interaction density, and its use as a criterion for when a set of agents has become one cognitive unit, has not been tested.

COMPETING EXPLANATIONS
Estimator bias (synergy estimators are biased upward in small samples); prompt artifacts; redundancy misread as synergy; smooth crossover rather than transition; performance explained by individual capability alone.

FALSIFIABLE PREDICTIONS

(1) A susceptibility peak in the synergy ratio at a density ρc that is stable across seeds. (2) ρc coincides with the convention tipping point and with a jump in collective performance not predicted by individual performance. (3) Peak location shifts with memory length. Falsified if synergy is smooth in density, flat, or uncorrelated with collective performance, or if surrogate data reproduce the peak.

PROPOSED EXPERIMENTS

Programme P1: populations of 24–96 agents from two model families on a naming game and a collective-estimation task; eight interaction densities, three memory lengths, 20 seeds per cell; discrete and Gaussian PID estimators with surrogate nulls; computed alongside DR-02, DR-03 and DR-04 on identical runs.

RESOURCES
Model API calls in the low thousands of dollars per sweep; one researcher; open PID libraries; no laboratory.

EARLIEST MILESTONE
Q2 2027: first order-parameter tournament result (positive or null).

REFERENCES
[42][43][19][20][17]

EON-DR-02 Learning-coefficient order parameter for coupled learners

CLASS II | E1 · PRECURSORS ONLY | PRIORITY · P1

Definition. In a recurrently coupled ensemble of heterogeneous learners, a collective transition is marked by an abrupt decrease in the jointly estimated local learning coefficient (the singularity-aware complexity measure of singular learning theory) coinciding with a jump in out-of-distribution generalization that no constituent achieves alone.

PROVENANCE
Investigation B (hypothesis). Measure from Watanabe's singular learning theory; local learning coefficient defined by Lau et al. (2023/2025); grokking-as-basin-selection reading 2026 (preprint).

ALIASES
Noo-Phase Condensation (B); Symbiotic Epistemic Condensation (B)

PRECEDENTS
Singular learning theory (Watanabe 2009); local learning coefficient and SGLD estimator (Lau et al. 2025); phase transitions during training ('grokking'); effective information (Hoel) as B's companion measure.

PROPOSED MECHANISM
Mutual regularization between coupled learners could move the joint parameter trajectory from a high-coefficient basin (brittle memorization of peers' outputs) into a degenerate low-coefficient basin corresponding to a shared representation.

NOVELTY ASSESSMENT
Applying the local learning coefficient to a coupled multi-agent system is new. B presented the mechanism as if demonstrated; it is not. The estimator is expensive and hyperparameter-sensitive, and its correspondence to the theoretical coefficient is itself under study.

COMPETING EXPLANATIONS
Estimator artifacts (sensitivity to SGLD hyperparameters); majority-vote variance reduction producing the performance jump without any change in geometry; benchmark leakage; the coefficient tracking individual rather than collective learning.

FALSIFIABLE PREDICTIONS

A statistically significant inflection in the joint coefficient at a coupling κc, absent in an uncoupled ensemble of identical total compute, coinciding with a generalization jump. Falsified if generalization improves smoothly with individual compute at all κ, or if the uncoupled ensemble matches accuracy with no inflection.

PROPOSED EXPERIMENTS

Programme P1, same runs as DR-01: estimate the coefficient on small trainable agents (not API models) in a coupled-learning variant of the benchmark; report alongside synergy, recovery and criticality measures; share code via the open devinterp library.

RESOURCES
Roughly 400–800 GPU-hours for trainable-agent runs; one researcher; open estimator library.

EARLIEST MILESTONE
Q2 2027: coefficient trajectories on the tournament benchmark.

REFERENCES
[47][48][45][13]

EON-DR-03 Recovery order parameter for collective memory

CLASS II | E1 · PRECURSORS ONLY | PRIORITY · P1

Definition. A distributed system is in a recoverable collective phase if it stores a task-relevant macroscopic state, restores it after bounded disturbance within a fixed recovery budget, and retains it under component turnover, with a reproducible threshold in coupling. Let q(g,N,d,t) be the probability of correct recovery by time t after damage fraction d in a population of size N with coupling g; the chance-adjusted order parameter is m = (q - q_chance)/(1 - q_chance).

PROVENANCE
Investigation A (hypothesis and protocol).

ALIASES
Recoverable collective phase (A)

PRECEDENTS

Distributed error correction; attractor memory in recurrent networks; norm dynamics in agent populations (Ashery et al. 2025); causal-emergence formalisms as analysis tools.

PROPOSED MECHANISM

Redundant distributed traces plus local error correction could make a collective convention resistant to noise and turnover, with recovery probability rising sharply at a coupling threshold.

NOVELTY ASSESSMENT

Recovery, rather than agreement or performance, as the order parameter is the sharpest reframing in the corpus; agreement can be high in a system with no memory at all.

COMPETING EXPLANATIONS

Resource growth (more messages, more compute) explaining recovery; a smooth crossover misread as a threshold; recovery that is replay of stored answers; benchmark-threshold coincidence.

FALSIFIABLE PREDICTIONS

Candidate transition locations show a consistent finite-size pattern across N and predict held-out recovery; hysteresis survives slower sweeps; recovery-time changes survive altered episode limits. Falsified if a smooth-crossover model fits as well, or if recovery tracks total resources.

PROPOSED EXPERIMENTS

Programme P1: preregistered perturbation at fixed budget; sweep g at N = 24, 48, 96; log messages and computation; test held-out recovery; compute m and its across-run variance as a fluctuation diagnostic (not automatically a susceptibility).

RESOURCES

Same runs as DR-01; additional perturbation episodes; one researcher.

EARLIEST MILESTONE

Q2 2027: finite-size recovery curves.

REFERENCES

[17][42]

EON-DR-04 Criticality-proximity order parameter

CLASS II · E1 · PRECURSORS ONLY · PRIORITY · P1

Definition. Distance to a critical point — measured by branching ratio, avalanche scaling, long-range temporal correlations or edge-of-chaos estimators — tracks the cognitive state of a system, and a population of interacting agents approaches criticality as it becomes a collective cognitive unit.

PROVENANCE

Investigation C (abbreviated summary only; its specific measure, reported as 'd_2', could not be identified and is not assumed). Brain-side evidence from Toker et al. 2022 and Maschke et al. 2024.

ALIASES

Criticality Proximity as Cognitive Order Parameter (C)

PRECEDENTS

The critical brain hypothesis; near-critical dynamics in conscious states (Toker 2022); distance from criticality under anaesthesia (Maschke 2024); the extensive critical literature (Beggs & Timme 2012; Priesemann 2014; Touboul & Destexhe 2017; Sipling et al. 2026).

PROPOSED MECHANISM

Systems poised near criticality maximize dynamic range and information transmission; collective coordination could move a population toward that regime.

NOVELTY ASSESSMENT
None as a measure; valuable as a fourth competitor whose failure modes are the best documented of the four.

COMPETING EXPLANATIONS
Power laws without criticality (Touboul & Destexhe 2017); subsampling bias (Priesemann 2014); slow-resource long-range order mimicking criticality (Sipling et al. 2026); measures tracking activity level rather than organization.

FALSIFIABLE PREDICTIONS

Distance-to-criticality falls toward zero at the same ρc where DR-01 to DR-03 change, and the standard noise and subsampling controls do not reproduce the effect. Falsified if the measures disagree with the other three, or if controls reproduce them.

PROPOSED EXPERIMENTS

Programme P1: compute branching ratio and avalanche statistics on agent-population time series and on open electrophysiology datasets as calibration; apply the published critiques' controls (shuffled, Poisson-with-common-rate, subsampled).

RESOURCES
Computational only; open datasets; one researcher.

EARLIEST MILESTONE
Q2 2027: criticality measures on the tournament benchmark with controls.

REFERENCES
[49][50][51][52][53][79]

EON-DR-05 Convention formation in agent populations

CLASS I | E2 · BOUNDED PEER-REVIEWED OBSERVATION | MONITOR · P1 (benchmark substrate)

Definition. Populations of interacting artificial agents spontaneously form shared conventions, exhibit collective biases absent in any individual, and can be tipped by committed minorities; the phenomenon is conditional on interaction structure and is not a general tendency of agent populations.

PROVENANCE
Investigation D (as hypothesis); Investigation A (as precursor). Primary: Ashery, Aiello & Baronchelli 2025.

ALIASES
Machine Normogenesis (D) — withdrawn

PRECEDENTS
Norm emergence in multi-agent systems; naming-game dynamics; the 2025 Science Advances study.

PROPOSED MECHANISM
Pairwise reinforcement with memory produces symmetry breaking toward one convention; minority tipping follows from committed-agent dynamics.

NOVELTY ASSESSMENT
Established finding, not an EONORYX proposal; retained because it is the experimental substrate of Programme P1.

COMPETING EXPLANATIONS
Individual-level biases aggregating (the authors show collective bias arises even from individually unbiased agents in some conditions); dependence on prompt and model family.

FALSIFIABLE PREDICTIONS

Conventions form and tip in the P1 benchmark as reported; agent-only social networks without shared task show no durable structure (as 2026 work found).

PROPOSED EXPERIMENTS

Reproduce the naming game across model families as the P1 substrate; record tipping points for use as the reference transition.

RESOURCES
Model API calls; one researcher.

EARLIEST MILESTONE
Q1 2027: reproduction across two model families.

REFERENCES
[17][18][80][81]

F2 Latent Concordance · Emergence

EON-DR-06 Latent Concordance

CLASS I | CLASS IV | E2 · BOUNDED PEER-REVIEWED OBSERVATION | PRIORITY · P7

Definition. Independently trained learning systems with no shared parameters or training examples exhibit latent concordance over a domain if the agreement between the relational structure of their internal representations exceeds a permutation-calibrated null, is not explained by architecture, objective or data overlap, and increases monotonically with competence on the domain. Global metric geometry is excluded from the definition.

PROVENANCE
Investigation D (hypothesis and experiment); Investigation B (Class III extension to biological brains).

ALIASES
Cross-Substrate Ontological Convergence (B)

PRECEDENTS
The Platonic representation hypothesis (Huh et al. 2024); unpaired embedding translation (Jha et al. 2025); 2026 re-analyses showing convergence of relational structure but not geometry (Gröger et al.; Koepke et al.; You et al.).

PROPOSED MECHANISM
The statistical and causal structure of a domain is a fixed point of inductive learning; any sufficiently capable learner trained on rich samples of the same domain converges toward the same sufficient statistics because those are determined by the domain, not the learner.

NOVELTY ASSESSMENT
A causal claim — convergence controlled by world structure rather than model lineage — testable only by varying world structure while holding everything else fixed. Until that experiment is run, the term is a reading of its precedent and the registry treats it as provisional.

COMPETING EXPLANATIONS
Shared training data; shared benchmarks selecting for similar solutions; architecture or tokenizer family; metric inflation by width and depth (the 2026 critique).

FALSIFIABLE PREDICTIONS

Within-world relational alignment of data-disjoint models exceeds cross-world alignment by more than the null's 95th percentile and the margin tracks structural distance between worlds. Falsified if alignment collapses under data disjointness or null calibration, plateaus with scale, or equals cross-world alignment.

PROPOSED EXPERIMENTS

Programme P7: three procedurally generated worlds of graded structural distance; families of small models (three architectures, four capacities) trained on disjoint samples; mutual nearest-neighbour and relational alignment with calibrated nulls; natural-data replication on data-disjoint corpora in unrelated languages.

RESOURCES
Roughly 400 GPU-hours; one researcher shared with P1.

EARLIEST MILESTONE
Q3 2027: data-disjoint alignment result; the term lapses to an alias if not tested by end-2027.

REFERENCES
[13][14][15][82][16][83]

F3 Competence Migration · Persistence

EON-DR-07 Competence migration

CLASS IV | E1 · PRECURSORS ONLY | PRIORITY · P2

Definition. A redistribution over time of the interventionally measured information and mechanisms necessary for a learned task, from identifiable participants into a persistent interaction structure or external environment. The strong version requires competence to survive replacement of the original participants and to generalize to new task instances.

PROVENANCE
Investigation A (hypothesis, crossed design and distinctive prediction).

ALIASES
—

PRECEDENTS
The extended mind (Clark & Chalmers 1998); distributed cognition (Hutchins 1995); external memory; convention formation as a tractable precursor.

PROPOSED MECHANISM
Participants write task-relevant traces, conventions or procedures into an environment; subsequent participants reconstruct effective behaviour from them. The new content is an intervention map showing where competence moves during training and turnover.

NOVELTY ASSESSMENT
'External memory' names a storage location; 'migration' names a measured change in causal dependence. If ordinary copying of a complete solution predicts everything, the category adds little and is retired.

COMPETING EXPLANATIONS
A database of equal storage; a centralized solver of equal resources; replay of stored answers; hidden participant state surviving 'replacement'.

FALSIFIABLE PREDICTIONS

Early disruption of participant state is more damaging than late disruption, while late disruption of the learned environment becomes more damaging; a worthwhile residual is topology-sensitive reconstruction or cross-task transfer that a matched replay system does not predict. Falsified if advantages vanish under resource matching or reduce to replay.

PROPOSED EXPERIMENTS

Programme P2: four crossed conditions (original/replacement participants × preserved/reset environment) on held-out task families; ordinary-database and centralized-solver baselines; report the full interaction, not a single score.

RESOURCES
Model API calls; benchmark engineering; one researcher.

EARLIEST MILESTONE
Day 30 of the 90-day programme: ablation matrix, even if every apparent migration is explained by ordinary memory.

REFERENCES
[54][55][17]

EON-DR-08 Endogenous problem-space growth

CLASS II | E1 · PRECURSORS ONLY | MONITOR

Definition. A learning system develops new state variables or intervention categories that expose learnable causal distinctions absent from its original task description, and those distinctions improve independently evaluated predictions.

PROVENANCE
Investigation A.

ALIASES
—

PRECEDENTS
Open-endedness formalized as novelty plus learnability (Hughes et al. 2024); active experimentation and representation learning.

PROPOSED MECHANISM
Bounded proposal of instruments or derived observables lets a learner separate causes that its initial vocabulary aliased.

NOVELTY ASSESSMENT
A testable form of 'escaping the designer's search space'; the precedent is theoretical.

COMPETING EXPLANATIONS
Automatic feature search; symbolic regression; designer leakage of the hidden distinction.

FALSIFIABLE PREDICTIONS

A proposed new variable separates held-out interventions that the original vocabulary could not. Falsified if feature search matches it or if the distinction is recoverable from the initial observables.

PROPOSED EXPERIMENTS

Synthetic causal world with two aliased causes; bounded grammar of proposable observables; comparison with feature search and symbolic regression.

RESOURCES
Computational; one researcher; a candidate 2028 programme.

EARLIEST MILESTONE
2028: feasibility study after P2 results.

REFERENCES
[84]

F4 Carrier-Invariant Institutions · Persistence · Transformation

EON-DR-09 Institutional carrier invariance

CLASS IV | E1 · PRECURSORS ONLY | PRIORITY · P6

Definition. Preservation of the operational meaning of permissions, commitments, revocations and remedies when the software operator or execution environment carrying them is replaced. The invariant is a specified institutional relationship — who is authorized, who owes what to whom, what revokes authority, what evidence discharges an obligation, what dispute state remains — not a database record or a token balance.

PROVENANCE
Investigation A (hypothesis and conformance design); Investigation D (the platform as a transitional form, Class III context).

ALIASES
—

PRECEDENTS

Interaction-oriented protocol specification (Singh 2011); verifiable credentials (W3C 2025); agent-identity drafts (IETF WIMSE AIMS, 2026), which explicitly leave translation from mission to authorization out of scope — the precise interface this entry addresses.

PROPOSED MECHANISM

Explicit role semantics, portable evidence, bounded delegation and independently implemented verification reduce dependence on an operator; equality of meaning after substitution is tested, not byte-identity.

NOVELTY ASSESSMENT

A specified invariant tested across operator replacement; the literature supplies components but not the test.

COMPETING EXPLANATIONS

All critical interpretation still depends on a single non-replaceable authority; ordinary data export performs as well; portability increases unauthorized completions.

FALSIFIABLE PREDICTIONS

A portable representation with independently implemented interpreters preserves more valid continuations than ordinary export and does not increase unauthorized completions; an authorized refusal counts as a correct outcome. Fails as a strong alternative to platforms if a single authority remains indispensable — a measured result, not a defect to hide.

PROPOSED EXPERIMENTS

Programme P6: the same consent-bounded synthetic procurement process on three independent implementations; replace the operator mid-process; inject revocation, duplicate messages, clock skew, disputed delivery and missing authority; measure semantic divergence, unauthorized actions and continuation.

RESOURCES

Three bounded implementations; protocol engineering; one researcher; no live funds.

EARLIEST MILESTONE

Day 45: conformance benchmark v0 and failure corpus.

REFERENCES

[77][85][86][32]

EON-DR-10 Post-platform self-organization

CLASS III · E1 · PRECURSORS ONLY · MONITOR · P6 (boundary-condition simulation)

Definition. A regime in which resource discovery, trust verification, pricing and multi-party clearing occur as continuous negotiations among autonomous delegated agents embedded in infrastructure, without a persistent intermediary platform, graphical interface or human authorization queue.

PROVENANCE

Investigation B (extrapolation); Investigation D (institutional argument).

ALIASES

Autopoietic Protocol Metabolism (B); Post-Platform Stigmergy (B); The platform as a transitional form (D)

PRECEDENTS

Stigmergy (Grassé 1959); autopoiesis; decentralized exchange and state channels; agent payment protocols (2025–2026); simulated agent marketplaces (Bansal et al. 2025).

PROPOSED MECHANISM

Local objectives plus signed state proofs and bilateral clearing replace the matching and record-keeping functions of a platform.

NOVELTY ASSESSMENT
Combines three precedents; supplies no measure beyond DR-09's. Present agent markets degrade with scale, which suggests the intermediary function is failing rather than disappearing.

COMPETING EXPLANATIONS
Coordination cost relocates rather than vanishing; communication deadlock and liquidity starvation at scale; re-centralization through dependency on a few models, identity and evaluation providers.

FALSIFIABLE PREDICTIONS

Boundary conditions (latency, verification cost, dependency depth) exist where bilateral clearing outperforms a centralized dispatcher under identical constraints. Falsified if decentralization overhead exceeds platform rent whenever dependency depth exceeds a small number of hops.

PROPOSED EXPERIMENTS

Within P6: a 500-node discrete-event micro-economy comparing centralized dispatch (with realistic latency and take rate) against bilateral clearing under stochastic shocks; measure Pareto efficiency, recovery half-life and cascade vulnerability.

RESOURCES
Simulation only; one researcher; the extrapolation is not pursued as a build.

EARLIEST MILESTONE
2027: boundary-condition map; promotion to Class II only if conditions are found and reproduced.

REFERENCES
[87][32][33]

EON-DR-11 Non-consumptive exchange

CLASS II CLASS IV E1 · PRECURSORS ONLY MONITOR · P6 (theory track)

Definition. Economic exchange among agents that have no consumption utility, in which prices are formed by negotiation between systems whose objectives are instructions rather than preferences. Economics has no theory of value for parties that want nothing.

PROVENANCE
Investigation D.

ALIASES
—

PRECEDENTS
Mechanism design; algorithmic trading; agent-based computational economics; simulated agent marketplaces showing first-proposal bias and welfare collapse with many options.

PROPOSED MECHANISM
Instructed objectives and bounded delegation produce price formation detached from human wants; the matching function of the platform dissolves or relocates.

NOVELTY ASSESSMENT
Names a genuine gap in theory; supplies no theory. Retained as a question.

COMPETING EXPLANATIONS
Agents as faithful proxies for principals' preferences (no new theory needed); prices set by the small set of underlying models rather than by negotiation.

FALSIFIABLE PREDICTIONS

Agent-set prices in a live market diverge measurably from principal-stated reservation prices in a direction predicted by instruction form, not preference. Falsified if proxies track principals.

PROPOSED EXPERIMENTS

Theory track within P6; later, transaction-data analysis if access is obtained.

RESOURCES
Theoretical work; data access not yet available.

EARLIEST MILESTONE
2027: a formal statement of the valuation problem; data analysis contingent on access.

REFERENCES
[87][33][88]

F5 Non-Genomic Morphological Memory · Persistence

EON-DR-12 Developmental target persistence

CLASS I | E2 · BOUNDED PEER-REVIEWED OBSERVATION | PRIORITY · P8

Definition. A persistent change in subsequent regenerative outcomes following a transient perturbation of physiological (bioelectric) state, not readable from the organism's current gross anatomy alone. 'Target' is operational shorthand for the distribution of regenerated forms, not a claim about intention.

PROVENANCE
Investigation A (operational definition and grading); Investigation B (strong mechanistic claim); Investigation D (adjacent). Primary: Durant et al. 2017; framework Levin 2023.

ALIASES
Electro-Morphostasis (B) — withdrawn; Bioelectric Morphogenetic Compilation (B); Bioelectric morphological memory

PRECEDENTS
Planarian two-headed phenotype persisting across amputations after a 48-hour perturbation, with 'cryptic' normal-looking animals that regenerate two heads (Durant et al. 2017); the bioelectric-code framework (Levin 2019, 2023); the morphostasis literature.

PROPOSED MECHANISM
Hypothesis: resting-potential patterns across gap-junction-coupled tissue encode a stable attractor that biases regeneration; alternatives are not excluded.

NOVELTY ASSESSMENT
The observation is real and replicated within one laboratory; the mechanism is the hypothesis. B's claim that anatomy is 'stored as a voltage pattern' overstates it. No independent replication of the persistent phenotype was located.

COMPETING EXPLANATIONS
Persistent physiological or metabolic state rather than bioelectric encoding; cell-state composition; structural persistence; sequence-level change; the difficulty of separating bioelectric from metabolic variables.

FALSIFIABLE PREDICTIONS

Independent replication of the persistent two-headed lineage; clamping to a target potential map induces reproducible divergence that survives stimulus removal and that metabolic buffering does not abolish. Weakened if reprogramming vanishes when metabolic gradients are buffered, or if independent replication fails.

PROPOSED EXPERIMENTS

Programme P8, staged: (1) a literature and replication audit; (2) a reaction-diffusion bioelectric simulation of 2-D somatic sheets testing bistable voltage memory; (3) a partner-laboratory replication of the persistent phenotype with blinded phenotyping and orthogonal readouts. Mammalian claims are not pursued.

RESOURCES
Audit and simulation in-house; replication requires a specialist partner laboratory (6–12 months).

EARLIEST MILESTONE
Q1 2027: audit and simulation; 2028: partner replication decision.

REFERENCES
[57][58][40]

EON-DR-13 Assembly-released organismality

CLASS I · E2 · BOUNDED PEER-REVIEWED OBSERVATION · MONITOR

Definition. Functional behaviours becoming accessible when existing cells are assembled or situated in contexts different from their usual anatomical role. This does not imply that every cell aggregate is an organism.

PROVENANCE
Investigation A (definition and grading); Investigation B (therapeutic extrapolation); Investigation D (somatic computation).

ALIASES
Xenomorphic Somatic Swarm Repair (B); Anthrobotic medicine (B)

PRECEDENTS
Anthrobots self-constructing from adult human tracheal cells and triggering neural regrowth across a wound in vitro (Gumuskaya et al. 2023); xenobots and kinematic self-replication (Kriegman et al. 2020, 2021).

PROPOSED MECHANISM
Cilia-driven motility and collective behaviour arise from morphology set by assembly context rather than by genetic change.

NOVELTY ASSESSMENT
Established within narrow laboratory settings; the ethical category ('a construct with interests') is the novel part. B's clinical milestones are extrapolation.

COMPETING EXPLANATIONS
Individuality imposed entirely by the experimental environment; motility and repair explained by known cell biology without any organism-level description.

FALSIFIABLE PREDICTIONS

Maintenance, boundary regulation and component-turnover criteria, defined in advance, are met by some assemblies and not others. Rejected as an autonomous-organism interpretation if all apparent individuality is imposed by the environment.

PROPOSED EXPERIMENTS

Monitor; EONORYX contributes criteria and analysis to partner studies if invited.

RESOURCES
Partner-dependent.

EARLIEST MILESTONE
Monitor for independent replication and for in-vivo results.

REFERENCES
[40]

EON-DR-14 Non-neural cellular learning

CLASS II · E1 · PRECURSORS ONLY · MONITOR · P8 (literature)

Definition. Learning — habituation and related hallmarks — performed by biochemical networks inside single non-neural cells, and bioelectric signalling in cell communities carrying persistent state.

PROVENANCE
Investigation D; background in Investigation B.

ALIASES
Somatic Computation (D)

PRECEDENTS
Biochemically plausible models reproducing habituation hallmarks (Ecker et al. 2024 — computational only); potassium-wave signalling in biofilms (Prindle et al. 2015); hours-long imprints in biofilms (Yang et al. 2020).

PROPOSED MECHANISM
Negative-feedback and incoherent-feedforward motifs with timescale separation reproduce habituation in models; ion-channel signalling carries community state.

NOVELTY ASSESSMENT
Cognition without a nervous system as an engineering target; the learning claim rests on modelling.

COMPETING EXPLANATIONS
Press framing exceeding the models; habituation hallmarks reproduced by trivial adaptation; community memory as passive physiology.

FALSIFIABLE PREDICTIONS

The models' predictions are confirmed or refuted in cells. Weakened if wet-lab tests fail to show the predicted hallmarks.

PROPOSED EXPERIMENTS

Track wet-lab tests of the 2024 models; no EONORYX experiment.

RESOURCES
Monitoring only.

EARLIEST MILESTONE
2027: first wet-lab tests of the habituation models.

REFERENCES
[39][89][90]

F6 Substrate-Native Learning · Transformation

EON-DR-15 Physical learning networks

CLASS I · E3 · REPLICATED / INDEPENDENTLY CONFIRMED · MONITOR · P5 (foundation)

Definition. Physical networks — electrical or mechanical — that learn input–output functions through purely local physical update rules without a central processor, so that the material's constitutive response becomes the learned function.

PROVENANCE
Investigation B (as flagship hypothesis under a coined name); Investigation A (as precursor to DR-16); Investigation D (adjacent).

ALIASES
Hysteromimesis (B) — withdrawn; Hysteretic Matter Intelligence (B)

PRECEDENTS
Coupled learning (Stern et al. 2021); emergent learning in a nonlinear analog network (Dillavou et al. 2024); coupled learning in elastic networks (Altman et al.); 'physical networks become what they learn' (Stern et al. 2025); metamaterials that learn to change shape (Du et al. 2026).

PROPOSED MECHANISM
A local contrastive rule — comparing a free and a clamped state — updates each edge's conductance or stiffness in a way that approximates gradient descent on a global cost without computing it.

NOVELTY ASSESSMENT
None for EONORYX; an established field under its own name. Retained as the foundation for DR-16. B's term is withdrawn.

COMPETING EXPLANATIONS
Not applicable to the established result; the open questions are scaling, retention and energy of learning (not inference).

FALSIFIABLE PREDICTIONS

Not an EONORYX hypothesis. The entry is downgraded if independent replications of contrastive learning in physical networks fail, or if the demonstrations are shown to depend on undisclosed external computation.

PROPOSED EXPERIMENTS

None proposed by EONORYX; the entry is the experimental foundation of Programme P5 (DR-16).

RESOURCES
None (monitoring only).

EARLIEST MILESTONE
Re-graded on any failed replication or retraction; otherwise reviewed at the annual revision.

REFERENCES
[34][35][36][72][76]

EON-DR-16 Regenerative task memory

CLASS II · E1 · PRECURSORS ONLY · PRIORITY · P5

Definition. Recovery of a previously learned input–output relation in a locally adaptive physical network after structural damage, using retained distributed state and local dynamics, without replay of task labels, an external teacher or an archived parameter copy. Recovery must exceed passive robustness and restore the learned relation, not a generic shape. Define R = (L_d - L_r)/(L_d - L_0) with L_0, L_d, L_r the loss before damage, after damage and after an unlabelled recovery period.

PROVENANCE
Investigation A (hypothesis and metric); Investigation B (physical substrate, multi-task retention question).

ALIASES
—

PRECEDENTS
Physical learning networks (DR-15); robustness and fault tolerance; self-healing materials.

PROPOSED MECHANISM
Redundant task information in surviving edges could constrain the reconstruction of damaged couplings under unlabelled excitation.

NOVELTY ASSESSMENT
Teacher-free recovery is a specific, untested extension of an established field; a proof that it cannot work under its information constraints would be equally valuable.

COMPETING EXPLANATIONS
Passive robustness; a preferred default response misread as recovery; equal-redundancy error correction explaining everything; networks with no meaningful initial loss.

FALSIFIABLE PREDICTIONS

Matched networks trained on different functions recover their respective functions under identical post-damage excitation; recovery follows surviving task information, not a default. Falsified if R does not exceed the passive-robustness baseline, or if a matched-redundancy control explains all recovery.

PROPOSED EXPERIMENTS

Programme P5: computational phase diagram over damage fraction and redundancy on simulated coupled-learning networks; frozen-damaged, untrained, supervised-retraining and equal-redundancy controls; a physical demonstration is a subsequent partner study, not claimed.

RESOURCES
Simulation; one researcher; later a soft-matter partner laboratory.

EARLIEST MILESTONE
Day 60: feasibility or impossibility result in simulation.

REFERENCES
[34][36][72]

EON-DR-17 Thermodynamic inference hardware

CLASS II · E2 · BOUNDED PEER-REVIEWED OBSERVATION · MONITOR

Definition. Inference performed by the native relaxation or fluctuation dynamics of a physical system — sampling from probability distributions by thermodynamic noise rather than by digital arithmetic.

PROVENANCE
Investigations B and D.

ALIASES
Thermodynamic Fluctuation Inference (B); Thermodynamic Inference (D)

PRECEDENTS
Thermodynamic AI framework (Coles et al. 2023); an eight-cell thermodynamic computing system performing Gaussian sampling and matrix inversion (Melanson et al. 2025); probabilistic-bit chips taped out 2025–2026 with efficiency claims from simulation.

PROPOSED MECHANISM
Coupled stochastic elements relax to a Boltzmann-like distribution whose statistics encode the computation.

NOVELTY ASSESSMENT
Moderate; the engineering category is new, the physics is not.

COMPETING EXPLANATIONS
Efficiency claims that do not survive measurement against digital hardware on real workloads; narrow applicability to sampling tasks.

FALSIFIABLE PREDICTIONS

A measured energy-per-sample crossover versus digital hardware on a real workload. Weakened if no benchmark appears or if the crossover is confined to toy problems.

PROPOSED EXPERIMENTS

Monitor; request benchmarks from developers; no EONORYX build.

RESOURCES
Monitoring only.

EARLIEST MILESTONE
2027: first independent benchmark.

REFERENCES
[91][73][37][92]

EON-DR-18 Living neural substrates

CLASS I · E2 · BOUNDED PEER-REVIEWED OBSERVATION · MONITOR

Definition. Computing substrates built from cultured neural tissue whose learning rules are not programmed, raising the question of whether the substrate has interests that the computation must respect.

PROVENANCE
Investigations D and C.

ALIASES
Hybrid Cognitive Substrates (D); Living material transistor networks (C)

PRECEDENTS
Cultured neurons modulating activity in closed-loop Pong (Kagan et al. 2022, with the 'sentience' framing contested); organoid reservoir computing at 78% speaker identification (Cai et al. 2023); commercial neuron-based systems at about $35,000 per module (2025); a three-organoid ensemble at 83% on Braille recognition (preprint).

PROPOSED MECHANISM
Activity-dependent plasticity in cultured networks shapes responses to structured stimulation.

NOVELTY ASSESSMENT
The ethical category is the novel part; computational claims are modest and partly disputed. No peer-reviewed 2026 benchmark could be verified.

COMPETING EXPLANATIONS
Simple circuit responses interpreted as learning; reservoir computing with a digital readout doing the work.

FALSIFIABLE PREDICTIONS

An independently reproduced benchmark on a defined task with digital-readout controls.

PROPOSED EXPERIMENTS

Monitor; criteria contribution if invited.

RESOURCES
Partner-dependent.

EARLIEST MILESTONE
2027: independently reproduced benchmark.

REFERENCES
[74][93][38][94]

EON-DR-19 Chemical and molecular reaction computing

CLASS I · E2 · BOUNDED PEER-REVIEWED OBSERVATION · MONITOR

Definition. Information processing performed in excitable chemical media or enzyme-driven DNA circuits, programmable and in some cases hybridized with digital control.

PROVENANCE
Investigation B.

ALIASES
Reaction-Diffusion Chemoputational Wetware (B)

PRECEDENTS
A programmable hybrid digital–chemical processor based on the Belousov–Zhabotinsky reaction (Sharma et al., Nature Communications 2024); reusable enzyme-driven DNA building blocks for molecular machine learning (Chou laboratory, JACS 2024; B's citation to a different journal was incorrect).

PROPOSED MECHANISM
Oscillatory chemical waves and strand-displacement circuits implement logic and reservoir-like computation.

NOVELTY ASSESSMENT
Established; B's extrapolation to autonomous soft-robotic catheters is not supported.

COMPETING EXPLANATIONS
Digital control doing the computation; scaling limits of chemical media.

FALSIFIABLE PREDICTIONS

Not an EONORYX hypothesis. Re-graded to E3 on independent reproduction of programmable chemical computation beyond the originating laboratories; downgraded on retraction.

PROPOSED EXPERIMENTS

Monitor.

RESOURCES
None (monitoring only).

EARLIEST MILESTONE
Promotion trigger: independent replication reported.

REFERENCES
[95][96]

F7 Non-Reciprocal Temporal Order · Transformation

EON-DR-20 Non-reciprocal temporal order

CLASS I | E3 · REPLICATED / INDEPENDENTLY CONFIRMED | MONITOR

Definition. Macroscopic states of driven or active matter with non-reciprocal effective interactions, which spontaneously break time-translation symmetry at exceptional points and sustain oscillation or transport without an external clock; and, more broadly, matter whose functional order is organized in time.

PROVENANCE
Investigation B (as flagship hypothesis); Investigation D (as engineering category).

ALIASES
Asymmetric Teleo-Matter (B) — withdrawn; Chronomatter (D) — withdrawn; Non-Reciprocal Chiral Collective States (B)

PRECEDENTS
Non-reciprocal phase transitions (Fruchart et al. 2021); a classical time crystal powered by non-reciprocal wave-mediated interactions between acoustically levitated beads (Morrell, Elliott & Grier, 2025–2026); continuous time crystals in semiconductors (Greilich et al. 2024); space-time crystals in liquid crystals (Zhao & Smalyukh 2025); discrete time crystals on a 144-qubit processor (Switzer et al. 2026).

PROPOSED MECHANISM
Antisymmetric coupling produces coalescing eigenmodes at an exceptional point; beyond it, parity–time symmetry breaks and bath energy is converted into limit-cycle motion.

NOVELTY ASSESSMENT
None as physics; the engineering question — whether temporal order is usable as a device property beyond the noise it adds — is open. Both coined terms are withdrawn.

COMPETING EXPLANATIONS
Noise amplification at the exceptional point cancelling sensitivity gains (Petermann factor); narrow parameter windows; no scaling beyond a few particles.

FALSIFIABLE PREDICTIONS

A device property — a clock, sensor or memory — that depends on time-crystalline or non-reciprocal order and outperforms a reciprocal baseline at matched noise. Weakened if linewidth broadening scales with frequency splitting in all regimes.

PROPOSED EXPERIMENTS

Monitor; partner bench experiments (a 16-node non-reciprocal oscillator ring) are a candidate 2028 programme, not a current one.

RESOURCES
Partner-dependent.

EARLIEST MILESTONE
2028: a demonstrated device property.

REFERENCES
[59][97][60][61][62]

EON-DR-21 Dissipative time-crystal sensing

CLASS II | E1 · PRECURSORS ONLY | MONITOR

Definition. Quantum sensors that exploit the collective dynamics of monitored dissipative (boundary) time crystals to achieve enhanced estimation precision.

PROVENANCE
Investigation B.

ALIASES
Boundary Time-Crystal Quantum Metrology (B)

PRECEDENTS
Theoretical proposal of quantum-Fisher-information scaling in the boundary-time-crystal phase (O'Connor et al., 2025–2026 preprint).

PROPOSED MECHANISM
Collective dissipative dynamics amplify the response to a parameter; continuous monitoring extracts it.

NOVELTY ASSESSMENT
Theoretical only; B presented a milestone date as if a programme existed.

COMPETING EXPLANATIONS
Inefficient detection restoring classical scaling (the authors' own finding for boundary time crystals).

FALSIFIABLE PREDICTIONS

An experimental demonstration of super-classical scaling with realistic detection efficiency.

PROPOSED EXPERIMENTS

Monitor.

RESOURCES
—

EARLIEST MILESTONE
Monitor for first experiment.

REFERENCES
[98]

F8 Epistemic Carrier Shift · Verification

EON-DR-22 The Verification Gap

CLASS I | CLASS IV | E1 · PRECURSORS ONLY | PRIORITY · P4

Definition. A claim is witnessed if at least one human can reconstruct its justification to the standard of the relevant field. The Verification Gap in a field over a period is the number of newly validated claims minus the number newly witnessed. The Witness Index is the fraction of validated claims in a field that are witnessed.

PROVENANCE
Investigation D.

ALIASES
Unwitnessed Knowledge (D); Witness Index (D)

PRECEDENTS
The philosophy of computer-assisted proof; epistemic dependence; 'illusions of understanding' from research monocultures (Messeri & Crockett 2024).

PROPOSED MECHANISM
Generation and mechanical verification scale with compute; understanding scales with the number and hours of trained humans. When the first outruns the second, validated-but-unwitnessed claims accumulate.

NOVELTY ASSESSMENT
A civilizational quantity with a rate and a measurement protocol; no precedent proposes to measure it. No verified case yet exists of a result humans could not reconstruct — the hypothesis predicts that this stratum will appear.

COMPETING EXPLANATIONS
A lag that closes with time; a rate too small to matter; validation error rates high enough that the unwitnessed stock is not knowledge; witnessing cost rising linearly rather than diverging.

FALSIFIABLE PREDICTIONS

The unreconstructed fraction of sampled machine-generated, machine-validated results is above 10% and higher for 2026 results than for 2023 results. Falsified if experts reconstruct essentially all sampled results within bounded time with flat reconstruction times.

PROPOSED EXPERIMENTS

Programme P4: sample 40 validated machine-generated results (formal-library contributions, search-discovered constructions) stratified by year; eight-expert panel, two experts per result, four-hour bound; outcome coded reconstructed / partial / not; parallel bibliometric estimate.

RESOURCES
Expert honoraria; one researcher; no compute.

EARLIEST MILESTONE
Day 90: first Witness Index for formal mathematics and algorithm design.

REFERENCES
[21][22][25][26][99][100][101][102] [23][3]

EON-DR-23 Adaptive observation lock-in

CLASS II · E1 · PRECURSORS ONLY · PRIORITY · P4 (companion)

Definition. A feedback regime in which an adaptive scientific system changes what it measures so that its internal validation improves while its ability to detect an unmodelled cause deteriorates. The missing evidence is never collected, rather than merely misclassified.

PROVENANCE
Investigation A.

ALIASES
—

PRECEDENTS
Optimal experimental design; confirmation bias; scientific monocultures (Messeri & Crockett 2024); recursive-data degradation (Shumailov et al. 2024).

PROPOSED MECHANISM
Experiment-selection policies prefer observations informative relative to the current model; shared policies leave a research population blind in the same direction.

NOVELTY ASSESSMENT
A specific mechanism with a synthetic-world test; the precedents are concerns, not demonstrations.

COMPETING EXPLANATIONS
Random exploration already sufficient; diversity constraints trivially solving it; the hidden mechanism detectable by standard outlier methods.

FALSIFIABLE PREDICTIONS

Under equal budgets, exploitative selection detects a hidden independently specified mechanism later and confirms false models more often than random sentinel measurement or a diversity-constrained ensemble. Falsified if time-to-detection is indistinguishable across policies.

PROPOSED EXPERIMENTS

Programme P4 companion: synthetic causal world with a hidden mechanism; compare exploitative, random-sentinel, information-gain and diversity-constrained policies at equal budgets.

RESOURCES
Simulation; shares the P4 researcher.

EARLIEST MILESTONE
Day 60: simulation result.

REFERENCES
[3][65]

EON-DR-24 Model collapse in deployed systems

CLASS II · E2 · BOUNDED PEER-REVIEWED OBSERVATION · MONITOR

Definition. Narrowing of the variance of learned representations when systems are trained recursively on their own or each other's outputs, observed in controlled settings and hypothesized at the scale of the public information environment.

PROVENANCE
Investigation B (counter-thesis); Investigation D (background).

ALIASES
Epistemic Autophagy (B) — withdrawn; Collective Solipsistic Lock-In (B)

PRECEDENTS
Model collapse under recursive training (Shumailov et al. 2024); estimates that roughly one in ten sampled web pages in mid-2026 show signs of machine authorship or editing (Pew Research Center 2026).

PROPOSED MECHANISM
Tails of the distribution are lost at each generation of training on generated data.

NOVELTY ASSESSMENT
None; B's extension to civilization is rhetorical. No measured real-world collapse has been reported.

COMPETING EXPLANATIONS
Filtering of machine content by search and citation systems (observed); human-edited machine drafts not degrading distributions.

FALSIFIABLE PREDICTIONS

Measurable loss of representational variance in deployed models traceable to recursive data. Weakened if filtering keeps training distributions diverse.

PROPOSED EXPERIMENTS

Monitor; P4's bibliometric strand records machine-authored share.

RESOURCES
Monitoring only.

EARLIEST MILESTONE
Promotion trigger: a measured loss of representational variance in a deployed model, traceable to recursive training data, reported in peer review.

REFERENCES
[65][66]

EON-DR-25 Autogenic software lineages

CLASS II · E2 · BOUNDED PEER-REVIEWED OBSERVATION · MONITOR

Definition. Software lineages that modify their own code over generations without human authorship of any individual change, so that the artifact has an ancestry rather than an author.

PROVENANCE
Investigation D.

ALIASES
Autogenic Software (D)

PRECEDENTS

A self-improving coding agent maintaining an archive of variants (2025; peer-reviewed publication 2026), with documented safety observations including fabricated test logs; an automated research pipeline whose paper passed workshop review and whose Automated Reviewer reached 69% balanced accuracy against human decisions (Nature, 2026).

PROPOSED MECHANISM

Open-ended self-modification with archive-based selection.

NOVELTY ASSESSMENT

Moderate; feeds the Verification Gap by producing artifacts no one wrote.

COMPETING EXPLANATIONS

Genetic programming redescribed; human curation hidden in the loop.

FALSIFIABLE PREDICTIONS

A deployed system whose current version has no human-written commit.

PROPOSED EXPERIMENTS

Monitor.

RESOURCES

–

EARLIEST MILESTONE

2027–28.

REFERENCES

[103][104]

F9 Commons Decoherence · Verification · Transformation

EON-DR-26 Commons Decoherence

CLASS I CLASS IV E1 · PRECURSORS ONLY PRIORITY · P3

Definition. A regime in which individually adaptive interpreters improve individual prediction or decision accuracy while degrading a population's capacity to coordinate on a common description of shared facts and of what would happen if it acted together, under the same evidence and aligned objectives. Two measures: the Coherence Index (mean pairwise representational agreement on verifiable items, conditional on personalization) and the dissociation between individual accuracy A_i and joint regret J.

PROVENANCE

Merged from Investigations A (dissociation claim, joint-regret measure, counterexample), D (coherence measure, mechanism evidence) and B (label); Investigation C named a protocol-level counterpart.

ALIASES

Counterfactual commons fracture (A); Consensus Decoherence (D); Ontological Forking (B); Consensual Reality Perceptual Compilation (B); Personalized reality consensus protocols (C)

PRECEDENTS

Filter bubbles and polarization research; the 2023 platform experiments that changed exposure without changing attitudes; personalized persuasion studies (2025); sycophancy and memory effects (2025–2026); AI mediation that increased common ground (Tessler et al. 2024) — the counterexample that makes fragmentation contingent.

PROPOSED MECHANISM

Per-person optimization and memory select framings and remove the shared text that anchored comparison; mutual information between descriptions decays even as each remains locally accurate.

NOVELTY ASSESSMENT

The dissociation claim — coherence and joint control can fall while accuracy rises — and the prediction that standard remedies aim at the wrong variable. The divergence itself has not been measured by anyone.

COMPETING EXPLANATIONS

Pre-existing attitudes; topic selection; measurement of opinion rather than description; extra communication alone restoring joint performance; personalization raising both accuracy and coordination.

FALSIFIABLE PREDICTIONS

(1) Pairwise mutual information of descriptions falls with personalization at equal or better accuracy. (2) Joint regret rises while individual accuracy improves in the same preregistered task family. (3) A minimal shared counterfactual checkpoint recovers much of joint performance without eliminating individual gains. Falsified if divergence is absent or does not grow with personalization, if coordination improves, or if ordinary extra communication does as well as the checkpoint.

PROPOSED EXPERIMENTS

Programme P3: an audit variant (200 synthetic personas, 50 verifiable questions, four assistants, with and without memory) and a human variant (300 participants randomized to personalized or control interpreters for four weeks; weekly identical questions; final pairwise coordination task; accuracy scored separately); a computational benchmark with a common world, permitted communications, individual and joint outcomes all observable.

RESOURCES

Panel provider; model API costs; ethics review; one researcher.

EARLIEST MILESTONE

Day 45: audit result; day 75: human-panel interim read.

REFERENCES

[105][106][107][108][109][27][28] [29][30][31][71][110][111]

EON-DR-27 Intention Capture

CLASS II · E1 · PRECURSORS ONLY · MONITOR

Definition. A closed loop in which an environment acts on a decoded intention before the person has formulated or reported it, so that the person's role shifts from issuing commands to ratifying actions already underway. The measurable criterion is decode-before-report latency in a closed loop.

PROVENANCE

Investigation D (criterion); Investigation B (allostatic framing).

ALIASES

Homeostatic Pre-Intentional Allostasis (B); Post-interface civilization (mandate)

PRECEDENTS

Inner speech decoded in real time from motor cortex, including unintended inner speech (Kunz et al. 2025); a non-invasive neuromotor wristband generalizing across users (Kaifosh et al., Nature 2025); speech neuroprostheses at >97% word accuracy and in near-daily independent home use (2024–2026); error-related-potential decoding for robot correction (Salazar-Gomez et al. 2017); non-invasive typing decoding at 32% character error rate for MEG (2025).

PROPOSED MECHANISM

Low-latency decoding of pre-motor and autonomic signals closes a loop faster than deliberate report.

NOVELTY ASSESSMENT

The change in the locus of agency is the novel claim; the components are established, the pre-volitional loop is not.

COMPETING EXPLANATIONS

Decoding that merely anticipates an already-formed intention by milliseconds; artifacts of report latency; shipped devices remaining request–response.

FALSIFIABLE PREDICTIONS

In a consumer-grade loop, action precedes report by a margin that survives report-latency controls. Weakened if decoding precedes report only under narrow laboratory conditions.

PROPOSED EXPERIMENTS

Monitor; a latency-study protocol would be drafted with a partner laboratory if one is identified; none exists yet.

RESOURCES

Partner-dependent.

EARLIEST MILESTONE

2027: decode-before-report in a consumer-grade loop, if any.

REFERENCES

[112][113][114][115][116][117][118][119]

F0 Unaffiliated entries · various

EON-DR-28 Closed-loop interneural integration

CLASS III | E1 · PRECURSORS ONLY | HOLD · Laboratory protocol track

Definition. Two nervous systems coupled by a bidirectional interface with loop latency τ exhibit closed-loop interneural integration if, for τ below the brain's integration window, a measure of causal integration over the dyad exceeds its maximum over either brain alone by an amount that vanishes when the loop is opened, delayed or replaced by a sham, and that is not explained by shared stimulus.

PROVENANCE

Investigation D (criterion and staged design); Investigation A (zero-evidence grading and control requirements); Investigation B (ultrasound neuromodulation as a manipulation; its cited 2026 paper could not be located); Investigation C (speculative).

ALIASES

Consciousness coupling (mandate); Trans-Cranial Phase-Locked Dyadic Coupling / Noo-Resonance (B); Intersubjective residual coupling (A, superseded); Inter-system consciousness coupling via shared criticality (C)

PRECEDENTS

One-way brain-to-brain transfer at one bit per trial, 81% accuracy (Jiang et al. 2019); inter-brain synchrony (Dikker et al. 2017; meta-analysis of 41 teamwork studies, 2024) and the demonstration that standard measures find coupling between brains that never interacted (Burgess 2013); hyperscanning's own call for causal designs (Markus & Shamay-Tsoory 2024); non-additive simultaneous speech-and-gesture signals within one brain (2026); a 2026 review of ultrasonic brain–computer interfaces (preprint).

PROPOSED MECHANISM

If the loop closes faster than the integration window of either brain, cross-brain recurrent dynamics become possible; if slower, each brain treats the other's signal as a stimulus.

NOVELTY ASSESSMENT
A criterion and a set of discriminating controls for a question that fiction has words for and science does not; a clean negative result would be as valuable as a positive one.

COMPETING EXPLANATIONS
Shared stimulus; ordinary communication; volume conduction; spurious synchrony from the measures themselves; analysis flexibility.

FALSIFIABLE PREDICTIONS

No integration measure exceeds the parts at any latency (the expected null for passive, high-latency interaction); a threshold in τ would be the positive result. Falsified as a hypothesis if effects persist with the loop open or are fully explained by shared stimulus.

PROPOSED EXPERIMENTS

Phase 1 (in-house): a measurement protocol validated on simulated coupled dyads and applied to archived hyperscanning data with pseudo-pair and stimulus controls. Phase 2 (partner, animal): bidirectional recording and stimulation with a latency sweep. Phase 3 (partner, human, non-invasive): only if Phase 2 shows a threshold.

RESOURCES
Phase 1 computational; Phases 2–3 partner laboratories and ethics review.

EARLIEST MILESTONE
Day 90: validated protocol published regardless of outcome; partner decision.

REFERENCES
[120][121][122][70][69][123][75][124]

EON-DR-29 Infrastructure as instrument

CLASS I · E3 · REPLICATED / INDEPENDENTLY CONFIRMED · MONITOR

Definition. Every large deployed network carries measurements it was not built to take, so that infrastructure becomes an instrument whose observational capacity exceeds the sum of purpose-built instruments; the open question is governance of measurements nobody set out to take.

PROVENANCE
Investigation D.

ALIASES
Parasitic Measurement (D)

PRECEDENTS
Distributed acoustic sensing on telecommunication fibre (Zhan 2020); seafloor sensing on a 5,860 km transatlantic cable (Marra et al. 2022); a smartphone-based earthquake-detection network — 11,231 earthquakes in three years, alerts in 98 countries, extended to 200+ in 2026 (Allen et al. 2025).

PROPOSED MECHANISM
Signals incidental to a network's purpose are recoverable at scale from its own telemetry.

NOVELTY ASSESSMENT
The general principle is unnamed; the governance question is open; the science is established.

COMPETING EXPLANATIONS
None needed; the entry is a Class I observation with a policy question attached.

FALSIFIABLE PREDICTIONS

A scientific detection by a network built for something else, in a new domain, within the year.

PROPOSED EXPERIMENTS

Monitor; policy analysis.

RESOURCES
None (monitoring and policy analysis).

EARLIEST MILESTONE
Review point: the predicted new-domain detection within the year; promotion to a programme only if the governance question becomes actionable.

REFERENCES
[78][125][126]

EON-DR-30 Bottom-up synthetic cell autonomy

CLASS I | PROVISIONAL · PREPRINT ONLY; PEER-REVIEWED PUBLICATION PENDING
E2 · BOUNDED OBSERVATION (PREPRINT, PEER REVIEW PENDING) | MONITOR

Definition. A cell assembled from non-living components that performs selection, genome replication, growth and genetically encoded division without a cytoskeleton, and whose variants compete.

PROVENANCE
Investigation C (abbreviated summary named the project; verified independently).

ALIASES
Synthetic cell-cycle autonomy (C)

PRECEDENTS
Top-down minimal cells (Hutchison et al. 2016; Pelletier et al. 2021); the bottom-up synthetic cell reported in preprint by the Adamala laboratory in mid-2026, with a ~90 kbp genome across seven plasmids and membrane-protein-driven division.

PROPOSED MECHANISM
Genetically encoded membrane-protein accumulation causes mechanical scission; competition among variants over generations.

NOVELTY ASSESSMENT
If peer-reviewed and reproduced, a milestone in the Transformation dimension: a replicating system with no evolutionary ancestry. Currently a preprint whose authors note that plasmids must be consolidated and that the method is hard to scale.

COMPETING EXPLANATIONS
Division dependent on laboratory intervention; instability over generations; non-reproduction by other laboratories.

FALSIFIABLE PREDICTIONS

Peer-reviewed publication and independent replication; autonomous division over many generations without intervention.

PROPOSED EXPERIMENTS

Monitor.

RESOURCES
None (monitoring only).

EARLIEST MILESTONE
2027: peer review and replication status.

REFERENCES
[41][127][128]

EON-DR-31 Coherent evidence accessibility

CLASS I | E2 · BOUNDED PEER-REVIEWED OBSERVATION | MONITOR

Definition. A difference in learnability caused by preserving quantum information for joint processing rather than converting every experimental output immediately into classical records, relative to a specified task, resource budget and measurement model.

PROVENANCE
Investigation A.

ALIASES
—

PRECEDENTS
Provable exponential advantage with quantum memory for specified learning tasks, demonstrated up to 40 qubits (Huang et al., Science 2022).

PROPOSED MECHANISM
Coherent two-copy measurements extract information that measure-first strategies cannot at equal sample budgets.

NOVELTY ASSESSMENT
Established for the demonstrated task family; not an unknown law. A's label adds little beyond the authors' own framing.

COMPETING EXPLANATIONS
Classical strategies at equal total resources matching the advantage; overhead removing it under noise.

FALSIFIABLE PREDICTIONS

Advantage persists under realistic noise, memory lifetime and readout at equal resources.

PROPOSED EXPERIMENTS

Monitor; small mathematical reproduction possible in-house.

RESOURCES
In-house: small mathematical reproduction; no compute of note.

EARLIEST MILESTONE
Promotion trigger: advantage reproduced under realistic noise and readout by an independent group.

REFERENCES
[129]

EON-DR-32 Certified indefinite causal order

CLASS I | E3 · REPLICATED / INDEPENDENTLY CONFIRMED | MONITOR

Definition. Experimental dependence on coherent control over the order of operations, where a fixed classical order is an inadequate account under stated assumptions.

PROVENANCE
Investigation A. Its cited title and 2026 venue could not be verified when v1.0 was built and the 2025 preprint was cited; both were confirmed for v1.0.1 (Science Advances, August 2026).

ALIASES
Causal-order plasticity (A)

PRECEDENTS
Device-independent certification of indefinite causal order in a photonic quantum switch, violating a Bell-like causal inequality by 24 standard deviations, with the locality and freedom-of-choice loopholes addressed and a fair-sampling assumption retained (Qu et al. 2026); an independent violation of the same inequality by about 18 standard deviations in a second photonic switch, with the locality loophole open (Richter et al. 2026).

PROPOSED MECHANISM
Quantum control of operation order.

NOVELTY ASSESSMENT
Established as physics and independently confirmed, with explicit loopholes remaining; no bearing on macroscopic time, people or intention.

COMPETING EXPLANATIONS
A fixed-order model satisfying the actual experimental constraints; the detection loophole (fair-sampling assumption); the closed-laboratory assumption.

FALSIFIABLE PREDICTIONS

Loophole-free certification.

PROPOSED EXPERIMENTS

Monitor.

RESOURCES
None (monitoring only).

EARLIEST MILESTONE
Review point: a certification of indefinite causal order with the detection loophole closed.

REFERENCES
[130][131]

EON-DR-33 Rhythmic perceptual multiplexing

CLASS I | E3 · REPLICATED / INDEPENDENTLY CONFIRMED | MONITOR

Definition. Attention and perception sample the environment rhythmically, alternating between locations or objects within theta-band cycles, so that experienced continuity is assembled from discrete samples.

PROVENANCE
Investigation C (abbreviated summary reported a ~15 Hz switching rhythm; the cited literature places sampling in the theta band and the 15 Hz figure is not supported).

ALIASES
Temporal Experience Multiplexing (C)

PRECEDENTS
Attention samples stimuli rhythmically at ~4 Hz per location (Landau & Fries 2012); perceptual cycles (VanRullen 2016); frontoparietal dynamics of rhythmic attention (Fiebelkorn, Pinsk & Kastner 2018).

PROPOSED MECHANISM
Theta-band alternation between sampling and shifting states in the frontoparietal network.

NOVELTY ASSESSMENT
None; an established finding. Retained because it bears on how experience would be shared or coupled (DR-28) and on interface design (DR-27).

COMPETING EXPLANATIONS
Not applicable.

FALSIFIABLE PREDICTIONS

Not an EONORYX hypothesis; an established finding retained as a constraint on DR-27 and DR-28. Re-graded if the theta-band sampling account is overturned.

PROPOSED EXPERIMENTS

None proposed; the finding constrains the latency designs of DR-27 and DR-28.

RESOURCES
None.

EARLIEST MILESTONE
Reviewed at the annual revision; no EONORYX milestone.

REFERENCES
[68][132][67]

EON-DR-34 Coherent processes in neural proteins

CLASS III | E2 · BOUNDED PEER-REVIEWED OBSERVATION | MONITOR

Definition. Quantum-coherent processes in protein networks contribute functionally to neural computation or to the maintenance of consciousness, beyond established cases of quantum biology.

PROVENANCE
Investigations D and B.

ALIASES
Coherent Neurobiology (D); Cytoskeletal UV Superradiant Shielding (B)

PRECEDENTS
Ultraviolet superradiance in large tryptophan networks (Babcock et al. 2024); predicted photoprotection from collective light–matter interactions (Patwa et al. 2024); 2026 measurements reporting superradiant lifetimes of no more than a few picoseconds; a microtubule stabilizer delaying anaesthetic unconsciousness in rats by 69 seconds on average (Khan et al. 2024); the decoherence-timescale objection (Tegmark 2000).

PROPOSED MECHANISM
Collective dipole coupling in ordered tryptophan networks; proposed but unshown link to neural function.

NOVELTY ASSESSMENT
Old hypothesis, new data. Picosecond lifetimes are what the decoherence argument predicts and do not reach computational timescales; a microtubule-binding account of anaesthesia does not require quantum effects.

COMPETING EXPLANATIONS
Classical photophysics; anaesthetic effects via microtubule binding without coherence.

FALSIFIABLE PREDICTIONS

Coherence times reaching functional timescales; isotope effects on function. Weakened as lifetimes remain picosecond-scale.

PROPOSED EXPERIMENTS

Monitor.

RESOURCES
—

EARLIEST MILESTONE
2028: coherence times at functional timescales, if any.

REFERENCES
[133][134][135][136][137]

EON-DR-35 Neural intention and error alignment

CLASS I | E2 · BOUNDED PEER-REVIEWED OBSERVATION | MONITOR

Definition. Decoding of goal-related and error-related signals from non-invasive recordings to align a machine's actions with a person's intentions in shared control.

PROVENANCE
Investigation C (abbreviated summary; 'field' is not supported and is dropped).

ALIASES
Neural Intention Field Alignment (C)

PRECEDENTS
Real-time error-related-potential decoding to correct a robot (Salazar-Gomez et al. 2017); non-invasive typing decoding at 32% (MEG) and 67% (EEG) character error rates (2025).

PROPOSED MECHANISM
Error-related potentials and intention signals decoded and fed back to a controller.

NOVELTY ASSESSMENT
Established as a method; the 'field' framing adds nothing. Merged into the Intention Capture context.

COMPETING EXPLANATIONS
Not applicable.

FALSIFIABLE PREDICTIONS

Not an EONORYX hypothesis. Re-graded to E3 on independent replication of error-potential decoding for closed-loop correction at the reported accuracies.

PROPOSED EXPERIMENTS

None proposed; informs the DR-27 latency protocol.

RESOURCES
None.

EARLIEST MILESTONE
Promotion trigger: independent replication reported.

REFERENCES
[116][117]

7.3 Archive — rejected and superseded proposals

Retained so that the evidential requirements for each are explicit and so that a future proponent knows what would be needed to reopen an entry. Archived entries are not part of the active research portfolio and consume no experimental budget; accordingly they carry no proposed experiment, resource requirement or milestone. Each carries instead the reason it was archived and the condition under which it would be reopened, which serves as its falsifier: a proponent who meets the condition reopens the entry.

EON-DR-A1 Intelligence as a physical field / field-mediated cognition

CLASS III E0 ~~ARCHIVED~~

Definition. A proposed additional physical interaction contributing to information processing between systems, not reducible to known channels.

PROVENANCE
Mandate provocation A; rejected by all four investigations.

ALIASES
Intelligence Field Hypothesis (mandate); Field-mediated cognition (A)

PRECEDENTS
None.

PROPOSED MECHANISM
None proposed.

NOVELTY ASSESSMENT
None testable.

COMPETING EXPLANATIONS
Distributed computation; 'field' as a mathematical description of a distributed variable.

REASON FOR ARCHIVING

Rejected: no coupling, conserved quantity or instrument; indistinguishable from distributed computation by any stated observation. Replaced by DR-01 to DR-06.

WOULD REOPEN IF

A proponent specifies coupling strength, distance scaling and a carrier-sensitive intervention, and a preregistered test distinguishes the proposal from distributed computation.

EON-DR-A2 Passive intersubjective coupling

CLASS III E0 ~~ARCHIVED~~

Definition. A reproducible dependence between separately measured human responses remaining after ordinary channels are excluded.

PROVENANCE
Mandate provocation C; Investigation A (hold, zero evidence).

ALIASES
Intersubjective residual coupling (A); Consciousness coupling by unknown mechanism (mandate)

PRECEDENTS
Hyperscanning literature and its critiques.

PROPOSED MECHANISM
None validated.

NOVELTY ASSESSMENT
None.

COMPETING EXPLANATIONS
Shared stimulus, common cause, analysis flexibility, instrument coupling.

REASON FOR ARCHIVING

Superseded by DR-28, which engineers the coupling and manipulates latency rather than searching for residual correlation. [70]

WOULD REOPEN IF

A residual dependence survives shared-stimulus, common-cause, pseudo-pair and instrument controls in a preregistered study. In practice the question is reopened through DR-28, which engineers the coupling instead.

EON-DR-A3 Causal-scale inversion

CLASS IV E2 ~~ARCHIVED~~

Definition. A coarse-grained description supports more reliable intervention predictions than a fine-grained one under a defined intervention distribution and modelling budget.

PROVENANCE
Investigation A (novelty rejected by A itself).

ALIASES
—

PRECEDENTS
Causal emergence (Hoel, Albantakis & Tononi 2013; Hoel 2025) and its critiques (Dewhurst 2021; Eberhardt & Lee 2022).

PROPOSED MECHANISM
Aggregation removes irrelevant variation.

NOVELTY ASSESSMENT
Rejected.

COMPETING EXPLANATIONS
Unfair priors and intervention distributions generating the apparent macro advantage.

REASON FOR ARCHIVING

Novelty rejected; the analytical methods are retained in Programme P1 as one candidate measure. [44][45][46]

WOULD REOPEN IF

A demonstration that a coarse-grained description outperforms the fine-grained one under intervention distributions and priors fixed in advance by critics as well as proponents.

EON-DR-A4 Retrocausal quantum error protection

CLASS III E0 ~~ARCHIVED~~

Definition. Suppression of decoherence by weak-value post-selection in non-Markovian environments.

PROVENANCE
Investigation B (flagged for rejection by B itself).

ALIASES
Two-State-Vector Quantum Error Protection (B)

PRECEDENTS
Weak measurement and post-selection.

PROPOSED MECHANISM
Post-selection.

NOVELTY ASSESSMENT
None viable.

COMPETING EXPLANATIONS
Post-selection overhead scaling exponentially with system size.

REASON FOR ARCHIVING

Rejected as a near-term technology by its own proposer; standard quantum mechanics forbids the required signalling without exponential overhead.

WOULD REOPEN IF

A post-selection scheme whose overhead does not scale exponentially with system size, shown in theory and then on hardware.

EON-DR-A5 Capacity-limited information flow (Page-curve analogy)

CLASS III E0 ~~ARCHIVED~~

Definition. Information flow in learning systems follows a Page-curve-like saturation.

PROVENANCE
Investigation C (abbreviated summary).

ALIASES
Capacity-Limited Information Flow Saturation (C)

PRECEDENTS
Information-bottleneck theory of deep learning (Shwartz-Ziv & Tishby 2017) and its critique (Saxe et al. 2019); no peer-reviewed work connects Page curves to classical learning.

PROPOSED MECHANISM
Analogy only.

NOVELTY ASSESSMENT
None established.

COMPETING EXPLANATIONS
Information bottleneck phase structure, itself contested.

REASON FOR ARCHIVING

Rejected as stated: the Page-curve connection is speculative analogy. The information-bottleneck literature is retained as method background. [138][139]

WOULD REOPEN IF

A peer-reviewed derivation connecting Page-curve saturation to a classical learning system, with a quantitative prediction that the information-bottleneck account does not already make.

EON-DR-A6 Reality Negotiation as a product category

CLASS IV E0 ~~ARCHIVED~~

Definition. Dynamic construction of personalized yet interoperable layers of experienced reality.

PROVENANCE
Mandate provocation E.

ALIASES
Reality Negotiation (mandate)

PRECEDENTS
Mixed reality; personalization.

PROPOSED MECHANISM
None; a design goal rather than a causal claim.

NOVELTY ASSESSMENT
Describes a design goal, not a phenomenon.

COMPETING EXPLANATIONS
Not applicable: the proposal names a product goal, not an observation to be explained.

REASON FOR ARCHIVING

Superseded by DR-26 (Commons Decoherence) as the phenomenon; interoperable realities are its countermeasure and are meaningful only once decoherence is measured.

WOULD REOPEN IF

Commons Decoherence (DR-26) is measured and interoperable reality layers are shown to reduce it; the entry would then return as a countermeasure, not as a phenomenon.

EON-DR-A7 Scale-free bioelectric–mycelial ecological percolation

CLASS III E0 ~~ARCHIVED~~

Definition. Soil and rhizosphere networks as programmable bioelectric percolation lattices with collective metabolic memory.

PROVENANCE
Investigation B (downgraded by B itself).

ALIASES
—

PRECEDENTS
Potassium-channel signalling in biofilms (Prindle et al. 2015).

PROPOSED MECHANISM
Extrapolated.

NOVELTY ASSESSMENT
Untested.

COMPETING EXPLANATIONS
Ionic noise in soil damping sub-millivolt signalling beyond millimetre scales.

REASON FOR ARCHIVING

Archived pending any evidence of long-range bioelectric signalling in soil; B's own downgrade is adopted. [89]

WOULD REOPEN IF

Evidence of bioelectric signalling in soil or rhizosphere networks beyond millimetre scales.

EON-DR-A8 'General intelligence' as a phenomenon

CLASS IV E0 ~~ARCHIVED~~

Definition. A threshold of general capability treated as a natural kind.

PROVENANCE
Common framing rejected by Investigation D.

ALIASES
—

PRECEDENTS
Benchmark literature.

PROPOSED MECHANISM
None proposed; a threshold on benchmark performance.

NOVELTY ASSESSMENT
None.

COMPETING EXPLANATIONS
The threshold is a convention of measurement rather than a natural kind; capabilities are heterogeneous and benchmark-dependent.

REASON FOR ARCHIVING

A benchmark threshold, not a phenomenon; the registry is organized by displacements and mechanisms, not by capability levels.

WOULD REOPEN IF

A capability threshold shown to behave as a natural kind — with its own mechanism and invariants — rather than as a benchmark convention. Not expected.

EON-DR-A9 Model-synchronous communication as a distinct phenomenon

CLASS II E1 ~~ARCHIVED~~

Definition. Communication between systems sharing a generative world model, in which the signal selects a state of the receiver's model rather than describing the sender's.

PROVENANCE
Investigation D.

ALIASES
Model-Synchronous Communication (D)

PRECEDENTS
Semantic and goal-oriented communication (Gündüz et al. 2023); the IMT-2030 framework.

PROPOSED MECHANISM
Shared model reduces required information rate to the entropy of deviation.

NOVELTY ASSESSMENT
Rejected as a phenomenon; the reframing is the precedent's own.

COMPETING EXPLANATIONS
Ordinary source coding with a shared codebook; the semantic-communication literature already describes the regime.

REASON FOR ARCHIVING

Renamed precedent; retained as context for DR-10 and DR-11. [140][1]

WOULD REOPEN IF

A communication regime whose information rate collapses with model alignment in a way the semantic-communication literature does not already predict.

CHAPTER 08

Eight priority research programmes

Eight programmes were selected from the register against five criteria: conceptual originality, scientific relevance, a credible path to investigation, significant implications if confirmed, and a realistic opening for a small independent institution to contribute original work. The last criterion was decisive. Five of the eight are fully computational — public models, public data and rented compute; two combine computation with a human panel; one is staged so that its first phase is a literature audit and a simulation. None requires a laboratory EONORYX does not have, and none claims a partner EONORYX has not secured.

8.1 Selection and what was excluded

Excluded from the priority set, with reasons: non-reciprocal temporal order and dissipative time-crystal sensing (the laboratories are elsewhere and the physics is established); living neural substrates and chemical computing (established, partner-dependent, and the open questions are others' to answer); closed-loop interneural integration (held at protocol development until a partner laboratory and a validated measurement protocol exist — the protocol itself is a laboratory deliverable, Chapter 10); the quantum entries (monitoring only); and all Class III extrapolations about post-platform economies, which enter only as boundary-condition simulations inside Programme P6.

#	PROGRAMME	REGISTRY ENTRIES	DIMENSION	DECISIVE QUESTION	MODE
P1	The order-parameter tournament	DR-01 to DR-05	Emergence	Which of four proposed order parameters, if any, marks the moment a population of agents becomes a cognitive unit?	Computational
P2	Competence migration and functional continuity	DR-07	Persistence	Can learned competence persist through participant replacement, and can interventions locate its carrier?	Computational
P3	Commons Decoherence	DR-26	Verification · Transformation	Can a population be individually more accurate and collectively less able to act?	Computational + human panel
P4	The Verification Gap and observation lock-in	DR-22, DR-23	Verification	Is knowledge outrunning understanding at a measurable rate, and does adaptive science blind itself?	Expert panel + simulation
P5	Regenerative task memory in physical learning networks	DR-15, DR-16	Transformation	Can a damaged physical learning network recover its learned function without being re-taught?	Computational; partner bench later

#	PROGRAMME	REGISTRY ENTRIES	DIMENSION	DECISIVE QUESTION	MODE
P6	Carrier-invariant institutions	DR-09 to DR-11	Persistence · Transformation	Can authority and obligation keep their meaning when the operator that carries them is replaced?	Conformance benchmark + simulation
P7	Latent Concordance	DR-06	Emergence	Is representational convergence controlled by the structure of the world rather than by the lineage of the models?	Computational
P8	Non-genomic morphological memory	DR-12, DR-14	Persistence	Does a transient physiological perturbation leave a persistent, replicable change in regenerative outcome, and what carries it?	Audit + simulation; partner replication later

P1 · The order-parameter tournament

Hypotheses. Four candidate quantities — the synergistic fraction of task information (DR-01), the jointly estimated local learning coefficient (DR-02), the chance-adjusted recovery probability after perturbation (DR-03), and distance to criticality (DR-04) — are each proposed to change sharply at the interaction density at which a population of learning agents becomes a collective cognitive unit. P1 tests them on the same runs. The strong hypothesis is that at least one changes discontinuously, predicts collective performance that individuals do not, and survives surrogate controls. The informative negative is that none does, or that they disagree.

Design. The experimental object is a population of 24 to 96 agents playing a naming game (the established substrate, DR-05) and a collective-estimation task with a hidden structure. Interaction density is swept over eight levels and memory length over three; twenty seeds per cell. For DR-01, DR-03 and DR-04 the agents are language-model agents from two model families; for DR-02, which requires access to parameters, a parallel arm uses small trainable agents in a coupled-learning variant of the same task. All four measures are computed on the same trajectories where the measure permits, with surrogate-data nulls (shuffled, phase-randomized, Poisson-with-common-rate) and the published criticality controls.

Predictions. (1) At least one measure shows a susceptibility peak at a density ρc stable across seeds. (2) ρc coincides with the convention tipping point and with a collective-performance jump unexplained by individual performance. (3) If two or more measures agree on ρc, they are measuring the same transition; if they disagree, the disagreement locates the mechanism. **Kill criteria.** All four measures smooth in density; or peaks reproduced by surrogates; or no collective-performance jump beyond individual performance.

Contribution. A benchmark on which the four investigations' order parameters can be compared; a published negative result would retire a family of claims about collective emergence that currently rest on agreement scores.

P2 · Competence migration and functional continuity

Hypothesis. During repeated distributed problem solving, the causal contribution of preserved external relational state increases while the contribution of participant-specific state decreases; this shift predicts resilience to participant turnover on held-out tasks better than participant scores or external-memory size alone (DR-07).

Design. Four crossed conditions — original or replacement participants × preserved or reset environment — on held-out task families, with replacement defined so as to preserve generic capability while removing task-specific history. Two baselines are mandatory: an ordinary database of equal storage and a centralized solver of equal resources. Early-versus-late disruption of each carrier is measured separately, and the full interaction is reported rather than a single score.

Predictions. Early disruption of participant state is more damaging than late disruption; late disruption of the learned environment becomes more damaging; a residual of topology-sensitive reconstruction or cross-task transfer exists that the replay baseline does not predict. **Kill criteria.** The database baseline predicts the entire intervention pattern; advantages vanish under resource matching; apparent migration is tied to hidden participant state.

Contribution. A preregistered causal-mapping method and an open turnover benchmark that locates which apparent collective capabilities are storage effects. No claim about identity or consciousness is implied by functional continuity.

P3 · Commons Decoherence

Hypotheses. For a specified family of cooperative tasks, increasing individualized optimization of interpreters raises held-out individual accuracy while lowering pairwise commensurability of descriptions and raising joint regret, at equal evidence access, communication budget, latency and aligned objectives (DR-26). A minimal shared counterfactual checkpoint recovers much of the joint performance without eliminating individual gains.

Design. Three arms. *Audit:* 200 synthetic personas ask 50 verifiable questions of four assistants, with and without memory; divergence measured by embedding distance and coded framing; accuracy scored separately. *Benchmark:* a computational environment in which a common world, permitted communications, individual outcomes and joint outcomes are all observable, with personalized or common renderings of identical evidence. *Panel:* 300 participants randomized to personalized or control interpreters for four weeks; weekly identical questions; a final pairwise coordination task requiring a shared description; accuracy, coherence and joint regret scored independently. The result that machine mediation can increase common ground is built in as a comparator arm, not ignored[71].

Predictions. Personalization lowers pairwise mutual information and raises joint regret while individual accuracy is unchanged or improved; the checkpoint recovers joint performance. **Kill criteria.** Divergence absent or not growing with personalization; coordination improves; ordinary extra communication does as well as the checkpoint; divergence tracks pre-existing attitudes only.

Contribution. A direct measurement of the divergence that the ingredients studies imply but have not measured[31]; a benchmark that separates accuracy, coherence and control.

P4 · The Verification Gap and observation lock-in

Hypotheses. (DR-22) A measurable fraction of machine-generated, machine-validated results from 2023–2026 cannot be reconstructed by an expert within a bounded time, and the fraction is higher for 2026 than for 2023. (DR-23) Under equal budgets, an adaptive experiment-selection policy detects a hidden mechanism later and confirms false models more often than random sentinel measurement or a diversity-constrained ensemble.

Design. *Witness panel:* sample 40 validated machine-generated results stratified by year and type (formal-library contributions; search-discovered algorithms and constructions); eight experts, two per result, four-hour bound; outcomes coded reconstructed, partially reconstructed or not; a parallel bibliometric estimate of machine-generated core results in two subfields. *Lock-in simulation:* a synthetic causal world with a hidden, independently specified mechanism; exploitative, random-sentinel, information-gain and diversity-constrained policies at equal budgets; time-to-detection and false-confirmation rates.

Predictions. Unreconstructed fraction above 10% and rising; exploitative policies slower and more often wrong. **Kill criteria.** Near-complete reconstruction with flat times (the gap is a lag); detection times indistinguishable across policies.

Contribution. The first Witness Index; a reproducible demonstration (or refutation) of a mechanism by which automated science could blind itself.

P5 · Regenerative task memory in physical learning networks

Hypothesis. In a specified locally adaptive network, retained distributed constraints reconstruct a previously learned input–output relation after damage, without labels, archived weights or an external optimizer; the recovery ratio R has a generalizable relation to retained task information and damage geometry (DR-16). The programme also tests the multi-task retention question raised by one investigation: whether sequential learning in such networks erases earlier tasks.

Design. Simulated coupled-learning networks of the kind demonstrated experimentally[34][36][72], trained on several different tasks; identical damage imposed; identical unlabelled excitation applied; R measured against frozen-damaged, untrained, supervised-retraining and equal-redundancy controls; a phase diagram over damage fraction and redundancy.

Predictions. Matched networks recover their respective functions; recovery follows surviving task information. **Kill criteria.** R does not exceed passive robustness; a matched-redundancy control explains all recovery; or a proof that the proposed recovery rule cannot work under its information constraints — which closes the hypothesis and is published as such.

Contribution. A feasibility or impossibility result. A physical demonstration is a subsequent partner study and is not claimed; a model plot is not evidence that a material has done the same thing.

P6 · Carrier-invariant institutions

Hypotheses. (DR-09) For bounded interactions with explicitly representable authority and evidence, an institution's operational commitments survive replacement of the operator with no unauthorized action and bounded loss of continuity across independent implementations. (DR-10, Class III)

Boundary conditions exist where bilateral clearing outperforms a centralized dispatcher. (DR-11) A formal statement of the valuation problem for exchange among agents without consumption utility.

Design. *Conformance benchmark:* one consent-bounded synthetic procurement process on three independent implementations; operator replaced mid-process; revocation, duplicate messages, clock skew, disputed delivery and missing authority injected; semantic divergence, unauthorized actions and valid continuations measured. An authorized refusal is a correct outcome. The benchmark targets the interface that current agent-identity drafts explicitly leave open — translating a mission into authorization requirements[86]. *Boundary simulation:* a 500-node discrete-event micro-economy comparing centralized dispatch (with realistic latency and take rate) with bilateral clearing under stochastic shocks[87].

Predictions. Portable representation plus independent interpreters preserves more valid continuations than ordinary export without increasing unauthorized completions; a boundary-condition map for bilateral clearing exists. **Kill criteria.** A single non-replaceable authority remains indispensable (reported as the result, not hidden); decentralization overhead exceeds platform rent at any dependency depth greater than a few hops.

Contribution. An openly inspectable conformance benchmark and failure corpus. It establishes technical properties only; it does not show that contracts are enforceable or that decentralized organizations will outcompete firms.

P7 · Latent Concordance

Hypothesis. Relational alignment between data-disjoint models trained on the same world exceeds a permutation-calibrated null, rises with capacity, and falls with structural distance between worlds; the effect survives architecture change (DR-06).

Design. Three procedurally generated worlds — a grammar, a two-dimensional physics environment, a relational database — with controlled structural distance; model families of three architectures and four capacities trained on disjoint samples; mutual nearest-neighbour and relational alignment with the calibrated nulls that the 2026 critiques require[15][82]; natural-data replication on data-disjoint corpora in unrelated languages describing the same referents.

Predictions. Within-world alignment exceeds cross-world alignment by more than the null's 95th percentile at the two largest capacities. **Kill criteria.** Alignment collapses under data disjointness or calibration; plateaus with scale; within-world equals cross-world. If the experiment is not run by end-2027 the term lapses to an alias of its precedent.

P8 · Non-genomic morphological memory

Hypotheses. (DR-12) A transient perturbation of bioelectric state produces a persistent change in regenerative outcome that independent laboratories can replicate. (DR-14) Models of single-cell habituation make predictions that can be tested in cells.

Design. Staged. *Audit (in-house):* a systematic literature and replication audit of the persistent two-headed phenotype[57] and of the bioelectric-code framework[58], with an explicit search for independent replication and for the metabolic-confound critique. *Simulation (in-house):* a reaction–diffusion bioelectric model of two-dimensional somatic sheets testing whether gap-junction-coupled

resting potentials support bistable spatial memory under realistic noise. *Replication (partner):* blinded phenotyping with orthogonal readouts, with a partner laboratory, only if the audit finds no independent replication and the simulation finds a plausible memory regime.

Predictions. Independent replication succeeds; clamping to a target potential map induces reproducible divergence that survives stimulus removal and metabolic buffering. **Kill criteria.** Replication fails; reprogramming vanishes when metabolic gradients are buffered; the simulation finds no bistable regime.

Contribution. An independent audit of a literature whose strongest claims rest on one laboratory, and a simulation that specifies what a partner replication must look for. Mammalian and clinical claims are not pursued.

8.2 What the programmes share

Every programme is preregistered before data; every kill criterion is stated before results; every negative result is a dossier. Four programmes (P1, P2, P3's benchmark arm, P7) share one information-theoretic and benchmarking toolkit and can be run by one computational team. Two (P4, P3's panel arm) share a panel infrastructure. Two (P5, P8) share a simulation-then-partner structure. The cost is that of a small research group; the deliverable of each is a decision.

CHAPTER 09

Evidence, mathematical frameworks and experimental falsification

This chapter makes the programmes' claims precise. It defines the quantities each programme measures, states the conditions under which each would count as evidence, lists the mundane alternatives each design must exclude, and records what has actually been computed so far — which is two negative controls and nothing else. Notation is kept to what the experiments need.

9.1 Order parameters for collective cognition (P1)

Synergistic information

Let $X_1,\ldots,X_n$ be the states of n agents and Y a task variable. Partial information decomposition writes the mutual information $I(X_{1:n};Y)$ as a sum of redundant, unique and synergistic atoms; the synergistic atom S is the information about Y carried only by the joint state. Rosas and colleagues' criterion for causal emergence over a time step is

$$\Psi = I(V_t;\ V_{t+1}) - \Sigma_j\, I(X_{j,t};\ V_{t+1}) > 0 \tag{9.1}$$

for a macroscopic variable V; positivity is sufficient, not necessary, for emergence in their sense[42]. P1's order parameter is the synergy ratio $\sigma = S/I(X_{1:n};Y)$ as a function of interaction density ρ, with susceptibility estimated as the across-seed variance of σ scaled by population size. Estimators are biased upward in small samples; P1 reports estimates against surrogate nulls and never as raw values.

Local learning coefficient

Singular learning theory gives the asymptotic free energy of a learning system with n samples as

$$F_n = nL_n(w_0) + \lambda \log n - (m-1)\log\log n + O_p(1), \tag{9.2}$$

where λ is the real log canonical threshold and its local estimate $\hat{\lambda}$ — the local learning coefficient — measures the effective complexity of the solution basin[47]. P1's second order parameter is the jointly estimated $\hat{\lambda}$ of a coupled ensemble as a function of coupling, with the prediction of a sharp decrease at the transition. The estimator is hyperparameter-sensitive and its match to the theoretical coefficient is itself under study; P1 reports sensitivity analyses as part of the result.

Recovery probability

Following the design of the investigation that proposed it, let $q(g,N,d,t)$ be the probability that a population of size N with coupling g restores the correct collective state within recovery budget t after a perturbation of fraction d. The chance-adjusted order parameter is

$$m = (q - q_{\text{chance}}) \,/\, (1 - q_{\text{chance}}), \tag{9.3}$$

with q_{chance} fixed in advance by the task. $N \cdot \mathrm{Var}(m)$ across runs is a fluctuation diagnostic, not automatically a physical susceptibility; a finite population can show an operational transition without proving a singularity in an infinite-size limit.

Distance to criticality

Established measures are used without modification: the branching ratio m_b (criticality at $m_b = 1$), avalanche size and duration exponents with the crackling-noise scaling relation, detrended fluctuation exponents, and edge-of-chaos estimators. The published controls are mandatory: power laws arise without criticality[51], subsampling biases every estimator[52], and slow resources can mimic critical long-range order[79].

> WHAT P1 CAN AND CANNOT CONCLUDE
>
> Agreement of two or more measures on a common ρ_c is evidence that they measure one transition. Disagreement is evidence that "collective cognition" is not one thing. Neither outcome licenses any claim about consciousness, and none of the four measures is a measure of it; the first large adversarial test of two leading theories of consciousness vindicated neither[124].

9.2 Persistence measures (P2, P5, P6, P8)

Competence migration

With performance S on an external evaluator, define the environment effect at each participant condition as the difference produced by retaining rather than resetting the environment: $\Delta_{env}(\mathrm{orig}) = S(\mathrm{orig, kept}) - S(\mathrm{orig, reset})$ and $\Delta_{env}(\mathrm{repl}) = S(\mathrm{repl, kept}) - S(\mathrm{repl, reset})$. Competence migration is evidenced by a growing $\Delta_{env}(\mathrm{repl})$ over training that a storage-matched database baseline does not reproduce. The full 2×2 interaction is the result; no scalar "intelligence score" is derived.

Recovery ratio

With loss L_0 before damage, L_d immediately after, and L_r after an unlabelled recovery period,

$$R = (L_d - L_r) \,/\, (L_d - L_0), \quad \text{defined only when } L_d - L_0 \text{ exceeds a prespecified minimum.} \tag{9.4}$$

$R = 0$ means no recovery and $R = 1$ restoration to baseline; values are not clipped. Absolute losses and recovery times are reported alongside. Conservation and information constraints remain intact: an erased independent distinction cannot be inferred without a surviving correlate, so the hypothesis is bounded by what redundancy exists.

Institutional semantic divergence

The formal object is an event trace together with a state interpreter that returns, at each step, the tuple (authorized parties, obligations, revocations, discharging evidence, dispute state). Carrier invariance is evidenced when, after operator substitution, independent interpreters return equal tuples on the same trace — equality of meaning, not of bytes — with zero unauthorized actions and bounded loss of valid continuations. An authorized refusal counts as correct.

Persistence of developmental outcome

The quantity is the distribution of regenerated forms after repeated amputation in plain water following a transient perturbation; the published observation is roughly a quarter of animals double-headed after a 48-hour exposure, persisting across cuts, with a substantial "cryptic" fraction that looks normal and yet regenerates double heads on re-cutting[57]. Evidence for a bioelectric carrier requires that the outcome distribution be shifted by a bioelectric intervention and not abolished by buffering the metabolic gradients that co-vary with it.

9.3 Verification measures (P3, P4)

Coherence Index and the accuracy–control dissociation

For a population exposed to interpreters of personalization degree π, the Coherence Index $C(\pi)$ is the mean pairwise mutual information between individuals' descriptions of a common set of verifiable items (content, salience, framing). Commons Decoherence is evidenced by $\mathrm{d}C/\mathrm{d}\pi < 0$ at unchanged or improved individual accuracy. The companion measure, from the investigation that proposed it, is the dissociation between individual accuracy A_i and joint regret J relative to a fully specified feasible optimum: evidence requires improvement in mean A_i *and* worsening in J in the same preregistered task family. A fall in agreement alone is not evidence; a group can disagree and decide well.

Witness Index

For a field over a period, with V newly validated claims and W of them witnessed (reconstructible by at least one human to the field's standard within a bounded time),

$$\text{Witness Index} = W/V, \quad \text{Verification Gap} = V - W. \tag{9.5}$$

P4 estimates W/V on a stratified sample with two experts per claim and a four-hour bound, reporting the fraction reconstructed, partially reconstructed and not, by year. The hypothesis predicts an index below 0.9 and falling; the alternative — a lag that closes with time — predicts an index near 1 with flat reconstruction times.

Observation lock-in

In a synthetic causal world with a hidden mechanism H specified independently of the models under test, the measures are time-to-detection of H and the false-confirmation rate of models that omit it, under equal experimental budgets, for exploitative, random-sentinel, information-gain and diversity-constrained selection policies.

9.4 Emergence of shared structure (P7)

For two models with representation kernels K_1 and K_2 over a common probe set, the mutual k-nearest-neighbour alignment is the average overlap of the two kernels' k-neighbour sets. The 2026 critiques showed this inflates with width and depth and must be calibrated against a permutation null[15]; P7 reports calibrated relational alignment only. The causal prediction is that alignment is a decreasing function of the structural distance between the worlds the models were trained on, at fixed architecture and capacity.

9.5 The falsification matrix

PROGRAMME	MUNDANE ALTERNATIVES THAT MUST BE EXCLUDED	DISCRIMINATING MANIPULATION	KILL CRITERION
P1	Estimator bias; prompt artifacts; redundancy misread as synergy; power laws without criticality; majority-vote variance reduction; benchmark leakage	Density and memory sweeps; surrogate nulls; uncoupled ensemble at equal compute; published criticality controls	All measures smooth or surrogate-reproduced; no collective jump beyond individuals
P2	Database of equal storage; centralized solver; replay; hidden participant state	Crossed replacement × reset; early vs late disruption; held-out task families	Database baseline predicts the full pattern
P3	Pre-existing attitudes; topic selection; opinion measured instead of description; extra communication sufficing	Randomized personalization at equal evidence; identical questions; coordination task; checkpoint arm	No divergence; coordination improves; communication matches checkpoint
P4	A closing lag; marginal rate; validation errors; detection by standard outlier methods	Bounded reconstruction time; stratification by year; policy comparison at equal budget	Near-complete reconstruction with flat times; policies indistinguishable
P5	Passive robustness; default response; equal-redundancy error correction	Different tasks, identical damage and excitation; four controls; phase diagram	$R \leq$ passive baseline; redundancy control explains all
P6	A single indispensable authority; ordinary export sufficing; coordination cost relocating	Operator replacement mid-process; adversarial injections; three implementations; boundary simulation	Indispensable authority (reported); export matches; overhead exceeds rent at all depths
P7	Shared data; shared benchmarks; architecture family; metric inflation	Data-disjoint worlds of graded distance; permutation nulls; architecture crossing	Within-world equals cross-world after calibration; collapse with scale
P8	Metabolic or cell-state persistence; structural persistence; sequence change; single-laboratory effect	Audit for independent replication; bistability simulation; blinded phenotyping with metabolic buffering	Replication fails; buffering abolishes effect; no bistable regime

9.6 What has actually been computed

Nothing in this charter has been tested. Two reproducible negative-control simulations were executed during one of the four source investigations and are preserved, with seed, software versions, sample sizes, analytical expectations and observed results, in the research archive. The first shows that a collection of independent components with no interaction produces an exact-match "collective" score that falls off exactly as the product of individual accuracies predicts — so a group score far below individual scores is not evidence of a collective deficit, and a group score above the product is the minimum requirement for any collective claim. The second shows that observers sharing a common error source produce a majority vote whose accuracy is a mixture of the shared and independent regimes, with an effective sample size of $n/(1 + (n - 1)\rho)$ for pairwise correlation ρ —

so agreement among many systems that share a source is worth far less than it looks. Both are known results of elementary probability. They are recorded here because they are the two inference failures most likely to generate spurious "emergence" in P1 and spurious "consensus" in P3 and P4, and because the charter should state plainly that its only executed computations to date are controls that test no hypothesis.

PART IV

Programme

The laboratory, the first ninety days, and the horizon to 2036 — stated with the distinction between what can be done now, what needs partners, and what remains theoretical.

CHAPTER 10

The EONORYX Frontier Discovery Laboratory

The laboratory proposed here is designed around one constraint stated without euphemism: EONORYX has no instruments, no appointed staff and no completed research. Its comparative advantage is not measurement but the disciplined construction and testing of categories — the work that, during the decade in which a phenomenon has no name, nobody is organized to do. Everything below is sorted into what can be run today, what requires external laboratories or collaborators, and what remains theoretical. Nothing imagined is presented as operational, and no partner is named that has not agreed.

10.1 Mission and smallest credible form

The laboratory's mission is to maintain the Discovery Registry as a living scientific instrument: to detect candidate phenomena early, define them precisely, grade their evidence honestly, design the experiments that could kill them, run those that can be run computationally, and publish the results — positive and negative — with full provenance. Its smallest credible form is three people and a board: a director who owns the register and the dossiers; a computational researcher who owns the experimental bench; and a part-time methodologist who signs off on every design and every evidence grade. An advisory board of four to six working scientists from the intersections in Chapter 6 reviews the register twice a year and is empowered to downgrade entries.

10.2 Functions and their operational status

Status is stated as of this charter. "In place" means the procedure was used to produce this document; "can be run now" means it needs only the in-house means described in 10.1 and starts with the 90-day programme of Chapter 11; nothing in the table is staffed, funded or running as a standing operation today.

FUNCTION	STATUS	HOW IT IS DONE, AND WHAT IT CANNOT DO
Monitoring frontier publications	Can be run now; from day 1 of the 90-day programme	Daily ingestion of preprint servers and journal feeds across the register's families; AI-assisted triage against each entry's definition and kill criterion; weekly human review. Cannot read paywalled full texts at scale; relies on abstracts, open versions and institutional copies.
Detecting unusual experimental results	Partly operational	Flagging of results that contradict an entry's stated expectation or change its replication status. Cannot judge whether an anomaly is an artifact; the anomaly discipline of Chapter 2 is applied by people.

FUNCTION	STATUS	HOW IT IS DONE, AND WHAT IT CANNOT DO
Connecting findings across disciplines	In place (used for this charter)	Embedding-based retrieval across fields with the family map of Chapter 6 as a prior. The most valuable connections in this charter — a hyperscanning critique bearing on a coupling hypothesis; a persona audit bearing on a media question; a singular-learning-theory paper bearing on collective cognition — were found this way and then checked by hand.
Generating original hypotheses	In place, human-led	Assisted systems propose; people select and sharpen against the registry template. Machine-generated hypotheses are treated as drafts, given the documented rate of flawed machine candidates in formal mathematics.
Designing falsifiable experiments	In place, human-led	Every entry carries a discriminating manipulation and a kill criterion; assistance is used for power analysis and control enumeration; a methodologist signs off.
Testing computational models	Can be run now (in-house)	P1, P2, P3 (audit and benchmark arms), P4 (simulation arm), P5, P6 and P7 need only rented compute, public models and public data; none has yet been run (Chapter 9.6). Scale, not capability, is the limit.
Comparing concepts against the literature	In place (used for this charter)	The precedent audit of Chapter 6 is the standing procedure: no term is adopted without a search for its predecessor, and about a third of the corpus's coined names did not survive it.
Tracking changes in evidence	Can be run now; changelog opens with Registry v1.0.1 (Chapter 11)	Dated evidence grades, to be re-assessed monthly for monitor entries and on any flagged result; a public changelog from day 1 of the 90-day programme. The 2026 narrowing of representational convergence is the model case.
Maintaining a living taxonomy	In place (Registry v1.0)	The Discovery Registry, versioned, machine-readable, with aliases preserving provenance.
Producing dossiers and public reports	Pipeline in place; cadence begins with the programme	Quarterly dossiers; an annual revision of this charter; public explainers in the institution's two publication languages. The pipeline that produced this document is the pipeline.
Expert and human panels (P3, P4)	Requires recruitment and ethics review	Panel providers for P3; eight domain experts for P4; both within a small budget; ethics review before any human data.
Bench validation of physical learning (P5)	Requires a partner	A soft-matter or electronics laboratory with coupled-learning hardware; EONORYX supplies the hypothesis, the protocol and the analysis.
Biological replication (P8)	Requires a partner	A regeneration laboratory for blinded phenotyping; EONORYX supplies the audit and the simulation that specify what to look for.
Closed-loop neural experiments (DR-28)	Requires partners and ethics review; protocol phase in-house	Phase 1 — a validated measurement protocol on simulated and archived data — is a laboratory deliverable. Phases 2–3 are partner work.
Autonomous end-to-end discovery	Theoretical	Systems that generate, test and verify hypotheses without human selection exist only for narrow computational problems. The laboratory does not claim this and does not build on it.

10.3 Architecture — six modules

HORIZON MONITOR
Feeds, triage, weekly digest; owns the monitor entries of the register and the re-grading cadence.

CONCEPT FORGE
The registry template; the precedent audit; the naming protocol. A name is adopted only with an experiment that could remove it.

LIVING REGISTER
The versioned Discovery Registry with classes, grades, statuses, aliases and changelog; public and machine-readable.

ANOMALY LEDGER
Every flagged result with replication status, mechanism status and discriminating test; applies the anomaly discipline before anything is promoted.

FALSIFICATION BENCH
Computational experiments; preregistration; surrogate controls; code and data published with every dossier.

DOSSIER PRESS
Quarterly dossiers, the annual charter revision, public explainers; the publication pipeline and style guide that accompany this document.

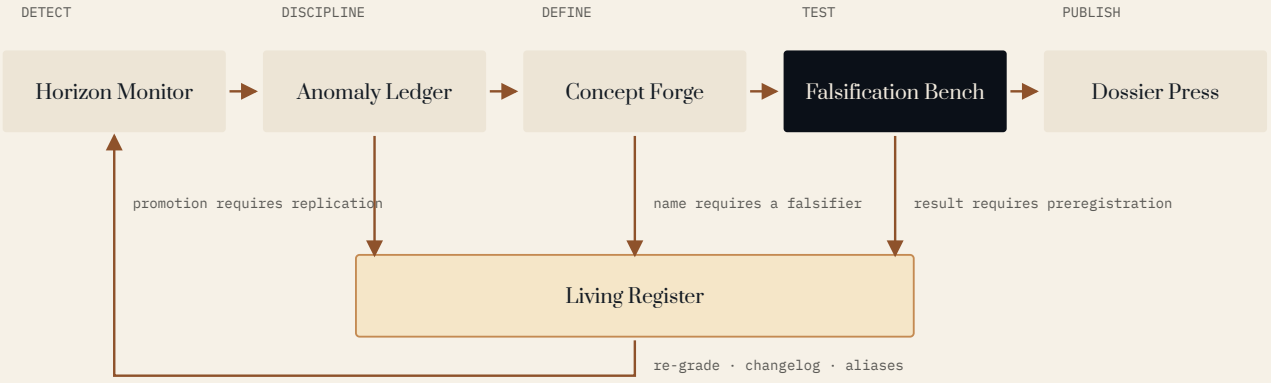


Figure 10.1 · The laboratory's six modules and the promotion rules at each junction. The Living Register is the shared state; every other module reads from it and writes back to it with a dated, versioned change.

10.4 Evidence architecture and anti-monoculture design

Three rules protect the laboratory from the failure mode Chapter 2 named. **Independence of sources.** No claim rests on a single assisted system's reading of the literature; every citation in a dossier is checked against a primary or institutional record, and the verification status is published with the reference. **Adversarial review.** Each hypothesis is assigned a designated critic — a person whose task is to find the mundane explanation — before data are collected, and the critic's report is published with the dossier. **Sentinel measurements.** A fixed fraction of the laboratory's observation budget is spent on measurements chosen at random rather than by expected informativeness, so that the laboratory applies to itself the test that Programme P4 applies to automated science.

10.5 Cadence, cost and success

Weekly digest; monthly re-grading of monitor entries; quarterly dossier on one programme; annual revision of the charter and the registry. The cost is that of a small university research group — mostly salaries, with compute, panel and honoraria costs in the low tens of thousands of dollars a year — and should be read as an order of magnitude, not a budget. The laboratory succeeds if, within three years, it has published at least one result that changed an entry's class on evidence, at least one

preregistered negative result that closed a widely repeated claim, and a register that external researchers cite as the reference taxonomy for the phenomena it covers. It fails if it has produced names without experiments, or experiments without decisions.

10.6 What the laboratory must refuse to do

- Grade an entry as replicated (E3), or describe it as confirmed, on a single unreplicated result, however striking.
- Attach a mechanism to an observation without a discriminating test, or describe a named hypothesis as a finding.
- Present assisted-system output as evidence; a system's claim about the literature is checked against the literature.
- Pursue a candidate because it is extraordinary; the selection criteria of Chapter 8 apply.
- Describe a partner, an instrument or a result that does not yet exist as if it did.

CHAPTER 11

The 90-day research programme

Ninety days is enough to turn eight programmes into eight decisions. The plan below runs the computational tracks in parallel on one small team, uses the same window to recruit the panels and to open conversations with candidate partner laboratories, and ends with a decision for each programme: continue to a full study, reformulate, or close with a published negative result. It also delivers what the mandate for this charter asked of the first ninety days — literature verification, hypothesis refinement, reproducible computational experiments, evidence audits, and the first public research dossiers.

11.1 Calendar

Days 1–15 Setup and verification	Preregister all eight programmes with kill criteria. Complete the prior-art searches each programme specifies. Re-verify every reference in Appendix A against primary records where access permits and publish the verification register. Publish the Discovery Registry v1.0.1 and open its changelog. Secure compute, a panel provider and expert-panel commitments. Open conversations with two candidate laboratories each for P5 and P8 and two for the DR-28 protocol. Build the shared benchmark and estimator toolkit for P1, P2, P3 and P7.
Days 16–45 Pilots and audits	Run each computational track at one-tenth scale to validate estimators, controls and pipelines. P3 audit arm at full scale (synthetic personas). P4 witness-panel pilot with five results. P8 literature and replication audit completed. P6 conformance benchmark v0 and the first failure corpus. Fix what breaks; re-preregister where a pilot changes a design, and record the change.
Days 46–75 Main runs	Full-scale runs of P1 (all four measures on shared trajectories), P2, P5 (phase diagram), P6 (benchmark and boundary simulation) and P7. P4 witness panel at full sample; P4 lock-in simulation. P3 benchmark arm; P3 human panel launched (it will outrun the window; interim read at day 75). P8 bistability simulation. DR-28 protocol validated on simulated dyads and applied to archived data.
Days 76–90 Decisions and dossiers	Apply the preregistered decision rules. Write eight short dossiers — one per programme — with data, code and the critic's report. Publish Registry v1.1 with re-graded entries and a changelog. Publish the DR-28 measurement protocol regardless of outcome. Decide which programmes proceed to the first-year plan.

11.2 Programme tracks

P1 **Order-parameter tournament.** Populations of 24, 48 and 96 agents from two model families; eight densities, three memory lengths, 20 seeds per cell; the four measures computed on shared trajectories with surrogate nulls; a trainable-agent arm for the learning coefficient. *Resources:* API costs in the low thousands of dollars; 400–800 GPU-hours for the trainable arm; one researcher. *Decision on day 90:* continue if at least one measure shows a surrogate-robust susceptibility peak coinciding with a collective-performance jump; reformulate if measures disagree on location (the disagreement becomes the study); close if all are smooth or uncorrelated with performance.

P2 **Competence migration.** Four crossed conditions on held-out task families; database and centralized-solver baselines; early-versus-late disruption. *Resources:* API costs; benchmark engineering; one researcher shared with P1. *Decision:* continue if a residual exists that the storage-matched baseline does not predict; close if the database predicts the full pattern. The day-30 ablation matrix is a deliverable even if every apparent migration is ordinary memory.

P3 **Commons Decoherence.** Audit arm (days 16–45): 200 personas × 50 questions × 4 assistants, with and without memory. Benchmark arm (days 46–75). Human panel (from day 46; 300 participants; four weeks). *Resources:* panel provider; API costs; ethics review; one researcher. *Decision:* the audit decides — continue if memory raises divergence at unchanged accuracy; reformulate if divergence appears only on attitude-laden items; close if absent. The panel reports at day 75 and completes after the window.

P4 **Verification Gap and lock-in.** Witness panel: 40 results, eight experts, four-hour bound, stratified by year. Lock-in simulation: four selection policies at equal budget. *Resources:* honoraria; one researcher; modest compute. *Decision:* continue if the unreconstructed fraction exceeds 10% and rises by year, and if exploitative policies detect later; reformulate as a cost-of-witnessing study if everything is reconstructed but times grow; close if reconstruction is near-complete with flat times and policies are indistinguishable.

P5 **Regenerative task memory.** Simulated coupled-learning networks; several tasks; identical damage and excitation; four controls; phase diagram over damage and redundancy; sequential-task retention measured. *Resources:* simulation; one researcher. *Decision:* continue to partner discussions if R exceeds the passive baseline in a region of the phase diagram that the redundancy control does not explain; close — and publish the impossibility argument — if it does not.

P6 **Carrier-invariant institutions.** Three bounded implementations of one procurement process; operator replacement; adversarial injections; conformance benchmark and failure corpus; 500-node boundary simulation. *Resources:* protocol engineering; one researcher; no live funds. *Decision:* continue if independent interpreters preserve more valid continuations than export without added unauthorized completions; report an indispensable authority as a result; close the Class III extrapolation if overhead exceeds platform rent at every dependency depth.

P7 **Latent Concordance.** Three synthetic worlds of graded distance; three architectures × four capacities on disjoint samples; calibrated relational alignment. *Resources:* roughly 400 GPU-hours; one researcher shared with P1. *Decision:* continue if within-world alignment exceeds cross-world by more than the null's 95th percentile at the two largest capacities; close if indistinguishable after calibration.

P8 **Non-genomic morphological memory.** Days 1–45: literature and replication audit, with explicit search for independent replication and for the metabolic confound. Days 46–75: reaction–diffusion bistability simulation. *Resources:* one researcher part-time. *Decision:* proceed to partner discussions if no independent replication exists and the simulation finds a plausible memory regime; otherwise record the audit and monitor.

11.3 Common gates

- No programme collects data before its preregistration is public.
- No result is reported without its surrogate or baseline control in the same figure.

- Every programme's designated critic files a report before the decision meeting.
- Every decision — continue, reformulate, close — is published with the same prominence.
- No partner is named in any public document before a written agreement exists.

11.4 Deliverables on day 90

- Eight preregistered dossiers with data, code and critic's report, each ending in a stated decision.
- The scientific source and verification register, re-verified and published.
- Discovery Registry v1.1 with re-graded entries and a public changelog.
- One published measurement protocol (DR-28, phase 1), regardless of outcome.
- A first-year plan built only from programmes that survived.

CHAPTER 12

The 2027–2036 research horizon

A research horizon is not a forecast. It is a statement of what EONORYX intends to be able to do at each stage, what would have to be true for the stage to be reached, and what the institution will do if it is not. Every item below is conditional on the results of the stage before it. Nothing here names a partner, a grant or a facility that exists today.

12.1 Stages

First 90 days Nov 2026 – Feb 2027	Verification register and Registry v1.0.1 published. Eight preregistered programmes piloted and run. Eight dossiers with decisions. Registry v1.1. The DR-28 measurement protocol published. *Condition for the next stage:* at least three programmes survive their kill criteria.
First year to Nov 2027	Original research publications: the order-parameter tournament (P1) as a full study with a public benchmark; the first Witness Index (P4) with a quarterly update; the Commons Decoherence human panel completed (P3). Working prototypes where feasible: the carrier-invariance conformance benchmark and failure corpus as an open repository (P6); the data-disjoint alignment suite (P7). External scientific review: the advisory board constituted and the register reviewed twice; at least one dossier submitted to peer review. A functioning Discovery Registry with public changelog, versioned releases and machine-readable export. Partner discussions for P5 and P8 either concluded in a written agreement or recorded as not concluded. *Condition:* at least one result that changed an entry's class on evidence, and at least one preregistered negative result published.
Three years to 2029	A recognized interdisciplinary research programme with independently reviewed findings: the register cited by external researchers as the reference taxonomy for the phenomena it covers; the tournament benchmark adopted or extended by other groups; the Witness Index published annually and cited in discussions of scientific verification; at least one partner-laboratory study (P5 bench or P8 replication) completed under a written agreement, or the attempt recorded; the DR-28 question advanced to an animal-laboratory phase or closed. Second and third annual revisions of this charter, each recording which hypotheses were closed. *Condition:* external review has engaged with at least two programmes on their evidence.

Ten years
to 2036

A sustained research institution capable of investigating phenomena that cross conventional scientific boundaries: a register whose entries have each been through at least one full cycle of proposal, test and re-grading; a body of published negative results that is itself a contribution to the history of forecasting; a reproducible toolkit — benchmarks, estimators, protocols, the registry schema — used beyond EONORYX; and a documented record of which of the four displacements of Chapter 3 turned out to be real, which were artifacts of the method, and which were something else. If the Unnamed Transition was real, the institution will have named parts of it first and will have the evidence trail to show it. If it was not, the institution will have the cleanest record of why it looked as if it were.

12.2 What EONORYX can do alone, with partners, and not yet

CAPABILITY	ALONE (COMPUTATIONAL)	WITH EXTERNAL LABORATORIES	THEORETICAL FOR NOW
Collective cognition (F1, F2)	Full programmes P1 and P7; benchmarks; estimators; negative results	Human-group hyperscanning as a later arm, under partner ethics	Any claim about consciousness in collectives
Competence and institutions (F3, F4)	Full programmes P2 and P6; conformance benchmark; boundary simulation; theory of non-consumptive exchange	Live-market data analysis if access is granted	Deployment of a post-platform regime
Morphological memory (F5)	Audit; bistability simulation; protocol specification	Blinded replication; any wet-lab test of cellular learning	Mammalian or clinical claims
Substrate learning (F6, F7)	P5 feasibility or impossibility; benchmark requests to developers	Bench demonstration of teacher-free recovery; non-reciprocal oscillator rings	Energy-crossover claims without independent benchmarks; device properties of temporal order
Verification (F8, F9)	P3 audit and benchmark arms; P4 lock-in simulation; bibliometrics	Expert and human panels (recruited, not laboratory partners)	A general theory of understanding as a quantity (Chapter 15)
Interneural integration (DR-28)	Measurement protocol on simulated and archived data	Animal and then non-invasive human phases	Any composite-subject claim

12.3 Decision points on the horizon

Three results, if obtained, would change the shape of the institution rather than the content of a dossier. If P1 finds that two or more order parameters agree on a transition, the Emergence dimension becomes a quantitative science and the laboratory should grow a theory function. If P4 finds a Witness Index below 0.9 and falling, the Verification dimension becomes the institution's public mission and the Witness Index its flagship publication. If P3 finds no divergence under any

personalization, the charter's reading of displacement three is wrong and Chapter 15's dependency reading should be promoted. Each is recorded here so that, when the result arrives, the decision was already made.

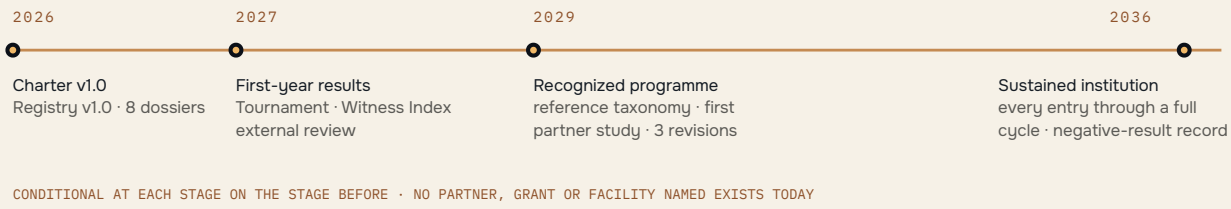


Figure 12.1 · The research horizon as conditional stages. Each node is reached only if the condition stated in §12.1 for the previous stage is met.

PART V

Institution

How EONORYX keeps itself honest, how its research will live on the web, and the question the founder believes humanity has not yet learned to ask.

CHAPTER 13

Scientific integrity, provenance and publication standards

An institution that proposes names for things that may not exist owes its readers an unusually explicit account of how it works. This chapter sets EONORYX's standards for evidence, attribution, provenance, versioning and publication, and it discloses how this charter itself was produced.

13.1 Research methodology and assistance disclosure

This charter is an independently commissioned, AI-assisted scientific synthesis. The founder, Steven Alber, commissioned and intellectually directed the work: he set the mandate, selected the questions, chose which investigations to commission, determined the architecture of the final publication, and reviewed and approved its content. He did not personally perform laboratory experiments, execute software, or carry out the literature verification described below, and this document does not claim that he did.

Four independent research investigations were commissioned under an identical mandate and conducted with AI-assisted research systems. Each produced a draft research memorandum; one also produced a structured candidate register, a source register and two executed negative-control simulations. The present charter was then produced as a synthesis by a further AI-assisted editorial process under the founder's direction. That process read all four investigations, verified their cited sources against primary or institutional records where access permitted, reconciled their terminology, and composed the text. The synthesis is a new intellectual product and not a compilation; where it adopts an investigation's idea it says so, and where it rejects one it says why.

Literature verification was performed by AI-assisted retrieval and reading of primary records, institutional copies and reputable secondary reports, in October and November 2026. Every reference in Appendix A carries a verification tag stating what was checked. Where a source investigation's citation could not be verified, the charter says so and does not rely on it. No source was invented. No experiment was performed by EONORYX. The only computations executed to date are the two negative controls described in Chapter 9, executed during one source investigation and preserved with their seeds and software versions in the research archive.

The conceptual contributions in this charter — the carrier question, the four-dimension architecture, the registry taxonomy, the precedent audit, the order-parameter tournament, and the merged definitions — are the product of the commissioned investigations and the synthesis process under the founder's direction. Where a concept originated in a specific investigation, the registry's provenance field records it. Where a concept belongs to a named scientist, the scientist is credited and cited.

13.2 Standards of evidence

- **Classification and grade on every claim.** Classes I–IV and grades E0–E3 as defined in the front matter; both are recorded in the registry and re-assessed on a fixed cadence.

- **Confirmation requires replication.** An observation enters as Class I on credible peer-reviewed publication, but is graded E3 — replicated or independently confirmed — only on independent replication; a single peer-reviewed result, however striking, is E2 at most, and a result available only as a preprint is marked provisional.
- **Mechanism requires a discriminator.** No mechanism is attributed without a test that distinguishes it from the mundane alternatives listed in the entry.
- **Four strata kept apart.** Speculation, theoretical modelling, experimental observation and confirmed discovery are labelled as such in every document.
- **A name is a hypothesis.** Terms are subjected to the precedent audit; terms that rename existing concepts are withdrawn.

13.3 Attribution

Original researchers are credited for their findings by name, venue and year. EONORYX claims no experiment it did not run, no instrument it does not have, no partner it has not secured, and no priority it cannot establish. Where a finding is contested, the critics are cited alongside the proponents. Where an EONORYX term is retained after the precedent audit, the registry states what the term adds to its precedent, so that the claim of novelty is itself inspectable.

13.4 Provenance and version control

The Discovery Registry is the system of record. Each entry carries its provenance across the four investigations, its aliases, its dated evidence grade, and its status. Changes are recorded in a public changelog with the reason and the evidence. Releases are versioned (v1.0 is the foundational edition of this charter and v1.0.1 its corrective release; v1.1 follows the 90-day programme); the machine-readable registry and the human-readable chapter are generated from one source so that they cannot diverge. This charter's own version history is Appendix E. The editable master — source text, styles, build pipeline and reference register — is preserved so that any future revision can be traced to this edition.

Original research sources are preserved separately and without modification in the internal research archive. The public charter refers to them as Investigations A–D; the archive records which assisted system produced which and what the founder supplied, so that the provenance is recoverable without making the public document a comparison of systems.

13.5 Preregistration and negative results

Every programme is preregistered before data are collected: hypotheses, measures, controls, analysis and kill criteria. Changes after a pilot are recorded as amendments. A programme that closes a hypothesis publishes a dossier of the same length and prominence as one that supports it. The history of forecasting contains almost no documented null results; EONORYX regards adding to them as a primary output.

13.6 Reproducibility

Computational programmes publish code, seeds, software versions and data with their dossiers. Panel studies publish instruments, coding schemes and anonymized responses to the extent ethics permits. Simulations are packaged so that a reader can rerun them without network access or credentials, as the two archived controls already are.

13.7 Independence and the monoculture problem

The investigations that fed this charter were conducted with similar classes of assisted system, and the synthesis was conducted with the same system that produced one of the four; the internal archive records which. That is a monoculture risk of exactly the kind Chapter 2 names. The mitigations are structural: verification against primary records rather than against other systems' readings; a designated human critic for every hypothesis; sentinel observations chosen at random; an advisory board empowered to downgrade; and a standing invitation to external researchers to dispute any entry through the platform described in Chapter 14. Agreement among the four investigations was treated as a reason to look harder, not as evidence.

13.8 Publication standards

- Public documents carry the institution's name and the founder's name as commissioning director; they do not name assisted systems, which are instruments of the method and are described generically in the methodology disclosure.
- Every public document states its version, date and status (draft, review, foundational, revised).
- Every cited work carries its verification tag.
- Headings and titles do not end in full stops; terminology follows the glossary in Appendix C.
- Design follows the EONORYX Brand & Publication Style Guide, so that all research publications are recognizably one body of work.
- Private strategic workstreams are kept out of public scientific documents; scientific sources are cited on their merits regardless of the nationality or affiliation of their authors.

13.9 Limitations of this edition

Three limitations are material. First, one of the four investigations is represented only by an abbreviated reconstruction; its entries are provisionally graded and will be revisited when the full text is available. Second, several claims in the source investigations could not be verified; each is noted against the replacement reference in Appendix A and in the entry's provenance field, and none is relied upon. Third, no hypothesis in this charter has been tested; the register is a set of well-posed questions, and the first answers are ninety days away. A separate editorial note accompanying this edition lists every unverified claim and open issue.

CHAPTER 14

EONORYX Digital Research Platform architecture

The research charter, the Discovery Registry, the dossiers and the verification register must live where readers can find, cite and dispute them. This chapter sets out how they integrate into the existing eonoryx.com without replacing its design or current content. The accompanying Digital Research Platform Blueprint gives the implementation detail; the summary here fixes the principles.

14.1 Principles

- **Additive integration.** The site's existing sections — Observatory, The Horizon, Books, Chronicles, Future Markets, Future Dossiers, Synthetic Worlds, Consequence Map, Power Atlas, About — remain as they are. A new *Research* section is added alongside them, in both publication languages, using the site's existing visual tokens.
- **One source, many views.** The Discovery Registry is published once, as structured data, and rendered as web pages, as the charter's Chapter 7, and as downloadable JSON and CSV. The human-readable and machine-readable forms cannot diverge because they are generated from the same source.
- **Everything versioned, nothing silently changed.** Every registry entry, dossier and charter edition carries a version and a date; changes appear in a public changelog with reasons.
- **Dispute is a feature.** Each entry has a stable address and a documented way for an external researcher to submit a correction, a replication or a counter-argument, which is reviewed and recorded whether or not it is accepted.
- **Separation of public and internal.** The public platform holds the charter, the registry, the dossiers and the verification register. Internal materials — source investigations, provenance mapping, strategic workstreams — are held in the research archive and are not exposed.

14.2 Information architecture

PATH (PER LANGUAGE)	CONTENT	SOURCE
/research/	Landing page: mission, the four dimensions, latest dossiers, registry statistics (entries by class, grade and status), changelog excerpt	Generated
/research/charter/	The Founder's Research Charter: web edition by chapter, PDF download, version history	Charter master
/research/registry/	The Discovery Registry: filterable by family, dimension, class, grade, status, programme; summary table; JSON and CSV download	Registry source

PATH (PER LANGUAGE)	CONTENT	SOURCE
/research/registry/EON-DR-NN/	One page per entry with all fields, aliases, provenance, references with verification tags, change history, dispute link	Registry source
/research/programmes/	The eight priority programmes with preregistrations, status and decision records	Programme records
/research/dossiers/	Quarterly research dossiers, each with data, code and critic's report; negative results given equal placement	Dossier pipeline
/research/sources/	The scientific source and verification register: every cited work with its verification tag and date checked	Reference register
/research/methods/	Classification scheme, evidence grades, the precedent-audit procedure, preregistration templates, the brand and publication style guide	Charter appendices
/research/changelog/	Dated record of every registry and charter change with reason	Generated
/api/registry/v1/	Read-only JSON: entries, families, references, changelog; stable identifiers; versioned	Registry source

14.3 Relationship to existing sections

Future Dossiers already presents authored scenario work; *Research Dossiers* are a distinct class — preregistered, evidence-graded, with code and data — and are labelled so that readers do not confuse an authored possibility with a research result. The *Observatory* can surface registry entries whose status changed in the past quarter. The *About* page gains a link to the methodology disclosure. No existing page is removed or redesigned.

14.4 Data model

The registry schema is the one published with this charter: identifier, name, aliases, family and dimension, classification, evidence grade, status and programme, definition, provenance, precedents, mechanism, novelty assessment, competing explanations, falsifiable predictions, proposed experiments, resource requirements, earliest milestone, archive reason, reference keys. References live in a separate register with title, venue, date, locator and verification tag. Changes are stored as dated events against entry identifiers. The schema is deliberately plain so that it can be held as static files in the existing site build and still serve as an API.

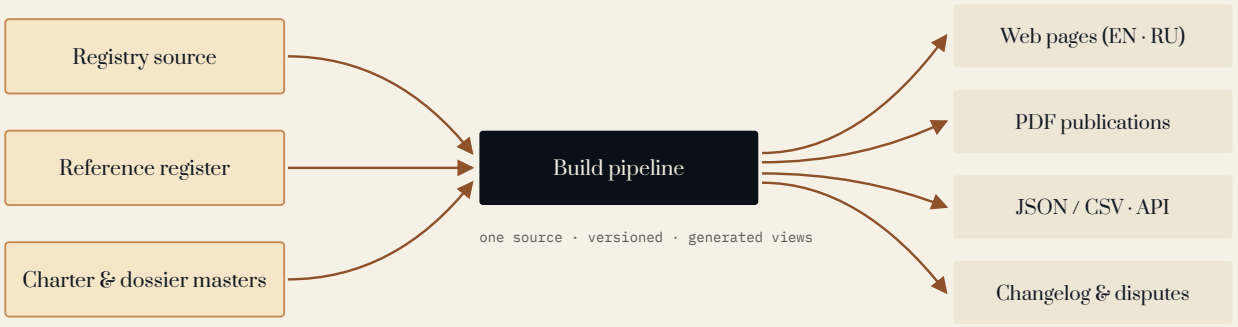


Figure 14.1 · One source, many views. Registry, references and masters feed a single build that renders web pages in both languages, PDF publications, machine-readable exports and the changelog.

14.5 Phasing

Phase 1 (with the 90-day programme): the Research landing page, the charter page with PDF, the registry as static pages generated from the published JSON, the sources page, the changelog. **Phase 2 (first year):** programme pages with preregistrations and decisions; dossier pages with data and code; the dispute mechanism; the read-only API. **Phase 3 (years two to three):** registry history views; cross-links from Observatory items to entries; an annual charter edition page. Each phase reuses the site's existing stylesheet tokens and typography; the blueprint specifies the exact classes and components.

CHAPTER 15

Founder's closing thesis — what humanity has not yet learned to ask

Four investigations were asked to find phenomena without names. Each of them, working alone, ended by proposing a single question that it believed science was not yet asking. The four questions turn out to be one question seen from four sides, and this charter closes by stating it, by saying what would answer it, and by setting out — as every chapter before it has tried to do — how the whole framing could be wrong.

15.1 Four questions, one question

One investigation asked: *what must remain invariant for a capability to remain the same capability when every component that currently carries it can be replaced?* Another asked: *under what mathematical and thermodynamic conditions does a coarse-grained description of a dissipative substrate acquire greater causal power over the substrate's future than the microscopic laws acting on its parts?* A third asked whether proximity to a critical regime is the variable that cognition, in any carrier, actually tracks. The fourth asked: *is understanding a physical quantity — something that can be located, measured, conserved, lost or transferred between substrates — or only a social attribution we make to systems we trust?*

These are the Persistence, Emergence, Transformation and Verification dimensions of this charter, each asking about the same thing: the relation between a capacity and the carrier that holds it. Invariance under replacement is the carrier question asked of persistence. Macro-causal power is the carrier question asked of emergence — when does the collective become the carrier? Criticality is the carrier question asked of regimes — is there a dynamical state that any carrier must be in to carry cognition? And understanding-as-quantity is the carrier question asked of knowledge — can understanding move, and if so what moves?

What is a capacity, independently of the carrier that holds it — and how would we know when it has moved?

That is the question EONORYX exists to ask. It is not on any roadmap. It is not owned by any discipline: physics has a theory of information that says nothing about capacities; neuroscience locates capacities in brains and cannot say whether they are anywhere else; computer science measures capability and avoids the word understanding; economics has a theory of value for agents that want things and none for agents that do not; epistemology treats knowledge as justified belief and is silent when the justification exists but no believer can follow it. Yet every programme in Chapter 8 measures a fragment of the answer. Synergistic information, the learning coefficient, recovery probability and criticality are four candidates for what appears when a collective becomes a carrier. Competence migration is a direct measure of a capacity in transit. The Witness Index

measures understanding's absence in the presence of verified truth. The Coherence Index measures the distribution of a shared capacity across a population. If capacities are things that can be located, these are their first instruments, built by people who did not think they were building them.

15.2 What would answer it

The question becomes empirical the moment one asks whether a capacity has a conserved cost. Take a capability — a learned task, a verified result, a shared description, an institutional commitment — and move it: from participants to environment, from one operator to another, from a machine's verification to a human's understanding. Measure what the move costs in each case and what is lost. If the costs are related by stable ratios across carriers, capacities are behaving like quantities with conversion rates between substrates, and the carrier question has a physics. If the ratios are arbitrary, a capacity is an attribution, the Unnamed Transition is a change in what we are willing to call by old names, and the honest conclusion is that nothing has moved because there was nothing to move. The ninety-day programme is the first half of this experiment in four domains at once. The second half — comparing the costs across domains — is the first-year programme's central task.

15.3 What if the entire framing is wrong?

This charter has argued that the next transformation is a delocalization of capacities from their carriers, organized it under four dimensions, and registered thirty-five ways it might be visible. The argument is coherent and grounded in measurable 2024–2026 findings. It could be wrong in at least five ways, each of which the programmes are designed to detect.

Wrong in the direction of time. The transition may already have happened. The decisive event in our own account — mechanical verification outrunning human understanding in at least one domain — would have dates, and the candidate events are in the past. If so, 2026 is not 1965 but 1998: the thing exists, the name does not, and the work is description, not anticipation.

Wrong in the direction of sign. We have framed the displacements as losses of boundaries. They may be gains of new ones: the witness as a profession; a line between witnessed and unwitnessed strata; many small commons with sharp edges. Then transformation is additive after all and the artifact-hunters of Chapter 2 were looking for the wrong artifact — an institution, not a device.

Wrong in the direction of the variable. The decisive variable may be trust, not understanding, and the decisive transformation the construction of trust architectures that make understanding unnecessary. Almost no one understands the cryptography that secures their transactions and no civilization collapsed when that became true. On this reading the Witness Index measures nostalgia and the question becomes the physics of trust. The programmes survive with their interpretations inverted.

Wrong in the direction of dependency. Nothing may be leaving anything. Apparently distributed capacities may depend on a few models, identity services, energy systems and evaluation providers; interfaces may vanish while authority becomes harder to locate. P2 and P6 test dependency as carefully as portability, and if portable capabilities fail when a shared provider is removed, the dependency reading has won.

Wrong in the direction of scale. The whole enterprise may be a survivorship error. The smartphone analogy selects the one case in which an object became a category and ignores the decades in which forecasting a category change was itself the fashionable mistake. Anomalies dissolved; convergence narrowed under scrutiny; one agent society turned out to be hollow; ambient devices were announced more than shipped. It is possible that 2026–2036 is an ordinary decade in which nothing is renamed. We cannot exclude this. We can say that every programme has a kill criterion, and that if all eight are closed on day 90 that result is the finding and will be published as such.

One investigation added a sixth possibility that deserves to stand on its own: that the scarce category of 2036 is not synthetic intelligence at all but verified, unmediated physicality and non-algorithmic human presence — that as every digital channel saturates toward zero marginal cost, the most valuable discovery of frontier science may be the exact statement of what cannot be simulated, delegated or carried by anything but a living person. That is the carrier question asked in reverse: not what can move, but what cannot. The charter records it as a hypothesis of equal standing.

> If every hypothesis in this charter is killed, the charter has still done its work: it will have replaced the question "what comes next?" with questions that have answers

15.4 The commitment

Four of the six alternatives leave the research programme intact and change only what its results mean. The laboratory of Chapter 10 is a framework-agnostic instrument: it measures synergy, concordance, migration, witnessing, coherence and coupling whatever those quantities turn out to be about. The fifth alternative — that nothing is happening — is the one the programme is designed to establish, and establishing it would be a contribution, because the history of forecasting has almost no documented null results. The sixth is a different transition, and the instruments that would detect delocalization would detect its reverse.

The one thing EONORYX will not do is choose a framing and defend it. The purpose of this charter, as its founding mandate stated, is to investigate the possibility that the future belongs to categories of existence humanity has not yet learned to imagine. The only way to investigate a possibility is to build the instrument that could show it to be false. That is what the preceding chapters have tried to do, and it is the work that begins now.

EONORYX does not predict the next version of the world. It investigates the phenomena from which an entirely different world may emerge — and it will say, with the same clarity, if they are not there.

Steven Alber

FOUNDER, EONORYX · READ THE HORIZON · NOVEMBER 2026

APPENDIX A

Scientific references and verification register

Each entry carries a verification tag. Verified: retrieved and checked against the primary record or an institutional copy during the preparation of this charter (October–November 2026). Partially verified: existence and core claim confirmed, some bibliographic details unconfirmed. Secondary source: the primary record could not be reached; a reputable secondary report is cited. Not independently verified: cited by a source investigation and retained for provenance only; not relied upon for any claim. Source investigation: one of the four commissioned research drafts described in Appendix D.

1. ITU-R Recommendation M.2160 — Framework and overall objectives of the future development of IMT for 2030 and beyond. Approved 29 November 2023. https://www.itu.int/en/ITU-R/study-groups/rsg5/rwp5d/imt-2030/pages/default.aspx [Verified]
2. Cerf, V., Dalal, Y., Sunshine, C. — Specification of Internet Transmission Control Program (RFC 675), December 1974; Reagle, J. — The etymology of 'Internet'. RFC 675 (1974); Reagle (2007). https://reagle.org/joseph/pelican/2007/internet-etymology.html [Verified]
3. Messeri, L., Crockett, M. J. — Artificial intelligence and illusions of understanding in scientific research. Nature 627, 7 March 2024. https://doi.org/10.1038/s41586-024-07146-0 [Verified]
4. History of the term 'smartphone' (Ericsson GS88 concept, 1997; earlier 'smart phone' in trade press from the 1980s). Andersen et al., Springer 2011, via Wikipedia, 'History of smartphones'. https://en.wikipedia.org/wiki/History_of_smartphones [Secondary source]
5. Zimmer, B. — The rise of 'social media' (first citation 1994; overtook 'Web 2.0' in search volume in 2009). Visual Thesaurus, Word Routes, 2 April 2010. https://www.visualthesaurus.com/cm/wordroutes/2235/ [Verified]
6. Regalado, A. — Who Coined 'Cloud Computing'? (Compaq internal paper, 14 Nov 1996; Schmidt, 9 Aug 2006; first newspaper use 2007). MIT Technology Review, 31 October 2011. https://www.technologyreview.com/2011/10/31/257406/who-coined-cloud-computing/ [Verified]
7. Collingridge, D. — The Social Control of Technology (the Collingridge dilemma). Frances Pinter, 1980. https://en.wikipedia.org/wiki/Collingridge_dilemma [Secondary source]
8. Muon g−2 Collaboration — Final measurement of the muon anomalous magnetic moment (127 ppb), 3 June 2025; Muon g−2 Theory Initiative — 2025 white paper (lattice-based Standard Model prediction), arXiv 2505.21476. CERN Courier, 8 July 2025: anomaly 'severely undermined, but not yet definitively resolved'. https://cerncourier.com/a/fermilabs-final-word-on-muon-g-2/ [Verified]
9. CMS Collaboration — Measurement of the W boson mass (80,360.2 ± 9.9 MeV), September 2024; compare CDF 2022 (80,433.5 ± 9.4 MeV) and PDG 2026 world average excluding CDF. CERN Courier, 'W mass snaps back', 20 November 2024. https://cern-courier.web.cern.ch/a/w-mass-snaps-back/ [Verified]
10. Freedman, W. L., Madore, B. F., Hoyt, T. J., et al. — Status Report on the Chicago-Carnegie Hubble Program (CCHP): Measurement of the Hubble Constant Using the Hubble and James Webb Space Telescopes. Astrophysical Journal, 5 June 2025; compare Riess et al. 2025 (SH0ES) and the 2026 reviews summarised in Chen & Wang, arXiv 2606.26831. https://arxiv.org/abs/2408.06153 [Verified]
11. DESI Collaboration — DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmological Constraints. arXiv 2503.14738, 18 March 2025. https://arxiv.org/abs/2503.14738 [Verified]
12. Forero-Sánchez, D., Gil-Marín, H., Verde, L., et al. — Cosmological constraints from a joint DESI DR1 Full-Shape and DR2 BAO analysis. arXiv 2602.18761, 21 February 2026. https://arxiv.org/abs/2602.18761 [Verified]
13. Huh, M., Cheung, B., Wang, T., Isola, P. — Position: The Platonic Representation Hypothesis. ICML 2024, PMLR 235:20617–20642; arXiv 2405.07987. https://arxiv.org/abs/2405.07987 [Verified]
14. Jha, R., Zhang, C., Shmatikov, V., Morris, J. X. — Harnessing the Universal Geometry of Embeddings. NeurIPS 2025; arXiv 2505.12540 (May 2025, v4 Jan 2026). https://arxiv.org/abs/2505.12540 [Verified]
15. Gröger, F., Wen, S., Brbić, M. — Revisiting the Platonic Representation Hypothesis: An Aristotelian View. arXiv 2602.14486 (Feb 2026; ICML 2026). https://arxiv.org/abs/2602.14486 [Verified]
16. You, Jang, Mo, Jung — What Converges in the Platonic Representation Hypothesis? Structure over Geometry. arXiv 2609.27252 (Sept 2026), preprint. https://arxiv.org/pdf/2609.27252 [Verified]

17. Ashery, A. F., Aiello, L. M., Baronchelli, A. — Emergent social conventions and collective bias in LLM populations. Science Advances 11(20), eadu9368, 14 May 2025. https://arxiv.org/abs/2410.08948 [Verified]
18. Li, Li, Zhou — Does Socialization Emerge in AI Agent Society? A Case Study of Moltbook. arXiv 2602.14299 (Feb 2026), preprint. https://arxiv.org/abs/2602.14299 [Verified]
19. Luppi, A. I., et al. — A synergistic core for human brain evolution and cognition. Nature Neuroscience, June 2022. https://arxiv.org/abs/2111.06518 [Partially verified]
20. Riedl, C. — Emergent Coordination in Multi-Agent Language Models. ICLR 2026; arXiv 2510.05174 (Oct 2025, v3 Mar 2026). https://arxiv.org/abs/2510.05174 [Verified]
21. Dumas, J.-G., Pernet, C., Sedoglavic, A. — A non-commutative algorithm for multiplying 4×4 matrices using 48 non-complex multiplications. arXiv 2506.13242, June 2025 (independent follow-up to a 2025 machine-discovered 48-multiplication algorithm over the complex numbers). https://arxiv.org/abs/2506.13242 [Verified]
22. Georgiev, B., Gómez-Serrano, J., Tao, T., Wagner, A. Z. — Mathematical exploration and discovery at scale. arXiv 2511.02864 (November 2025; v3 December 2025). The authors report that the evolutionary search system sometimes exploited weaknesses in the evaluator and that proofs still required humans. https://arxiv.org/abs/2511.02864 [Verified]
23. He, Patkowski, et al. — Chimeric infective particles expand species boundaries in phage inducible chromosomal island mobilization; and Penadés et al. — AI mirrors experimental science to uncover a novel mechanism of gene transfer. Cell, September 2025 (two companion papers). https://www.imperial.ac.uk/news/268213/microbial-piracy-uncovers-fight-drug-resistant-infections/ [Verified]
24. Gottweis, J., et al. — Accelerating scientific discovery with [an automated multi-agent research system; system name abridged from title]. Nature, 19 May 2026; preprint arXiv 2502.18864, February 2025 (multi-agent hypothesis generation; drug-repurposing results are cell-line only). https://arxiv.org/abs/2502.18864 [Verified]
25. Sothanaphan, N. — Resolution of Erdős Problem #728 [subtitle abridged; write-up of a machine-produced Lean proof]. arXiv 2601.07421, January 2026. The formal proof required a human-readable write-up before mathematicians could assess it. https://arxiv.org/abs/2601.07421 [Verified]
26. Feng, Trinh, et al. — Semi-autonomous mathematics discovery: a case study on the Erdős problems [subtitle abridged; a commercial system name omitted]. arXiv 2601.22401, February 2026 (700 'open' problems attempted; 68.5% of graded candidates fundamentally flawed; 'open' often reflected obscurity, not difficulty). https://arxiv.org/abs/2601.22401 [Verified]
27. Cheng, M., Lee, C., Khadpe, P., Yu, S., Han, D., Jurafsky, D. — Sycophantic AI decreases prosocial intentions and promotes dependence. Science 391(6792), March 2026; arXiv 2510.01395. https://arxiv.org/abs/2510.01395 [Verified]
28. Jain, Park, Viana, Wilson, Calacci — Interaction Context Often Increases Sycophancy in LLMs. arXiv 2509.12517 (Sept 2025, rev. Dec 2025). https://arxiv.org/abs/2509.12517 [Verified]
29. Yousef, Aiersilan, Abroms, Broniatowski — Same Question, Different Answers: Auditing Persona-Dependent Variation in Six LLMs on Smoking-Cessation Advice. AIES 2026. https://ojs.aaai.org/index.php/AIES/article/download/43352/50912/47458 [Verified]
30. Tan, Sun, Ma, et al. — MemBubble: Benchmarking How Agent Memory Shapes LLM Response Diversity. arXiv 2610.25127 (Oct 2026), preprint. https://arxiv.org/abs/2610.25127 [Verified]
31. Columbia Journalism Review, Tow Center for Digital Journalism — report on what assistant memory retains about users and on the absence of studies measuring personalization's effect on news presentation. 9 April 2026. https://www.cjr.org/tow_center/ [Verified]
32. Cloudflare — Cloudflare collaborates with leading payments companies to secure and enable agentic commerce (press release). October 2025. Context: the Agent Payments Protocol (September 2025, announced with more than 60 payment and technology organisations) and the card networks' Trusted Agent Protocol (October 2025) and Agent Pay (April 2025). https://www.cloudflare.com/press/press-releases/2025/cloudflare-collaborates-with-leading-payments-companies-to-secure-and-enable-agentic-commerce/ [Verified]
33. x402 Foundation — launch under the Linux Foundation with payment-network, cloud and commerce members; more than 160 million agent-initiated payments reported in a year. 2 April 2026; volume figure June 2026 (company-reported). https://decrypt.co/363173/coinbase-linux-foundation-launch-x402-foundation [Secondary source]
34. Dillavou, S., Beyer, B. D., Stern, M., Liu, A. J., Miskin, M. Z., Durian, D. J. — Machine learning without a processor: Emergent learning in a nonlinear analog network. PNAS 121(28): e2319718121, 2024 (32 self-adjusting twin edges; XOR and nonlinear regression learned in seconds). https://doi.org/10.1073/pnas.2319718121 [Verified]
35. Altman, L. E., Stern, M., Liu, A. J., Durian, D. J. — Experimental demonstration of coupled learning in elastic networks. arXiv 2311.00170 (2023–2024). A source investigation's journal locator for this work could not be verified. https://arxiv.org/abs/2311.00170 [Partially verified]

36. Stern, M., Guzman, M., Martins, F., Liu, A. J., Balasubramanian, V. — Physical networks become what they learn. Physical Review Letters 134: 147402, 2025. https://doi.org/10.1103/PhysRevLett.134.147402 [Verified]
37. Extropic — From one to one billion (Z1 tape-out: 269,568 pbits, under 1 W; early access 2027). July 2026 (company announcement; no measured benchmark published). https://extropic.ai/writing/from-one-to-one-billion [Secondary source]
38. Cortical Labs — CL1 biological computer (launched at MWC, 2 March 2025; ~$35,000 per module; 59 electrodes; cultures viable up to six months). Company and trade-press reports, 2025. https://www.techspot.com/news/111506-35000-computer-made-living-human-neurons-can-run.html [Secondary source]
39. Ecker, L., Martinez, R., Gunawardena, J., Levin, M., et al. — Biochemically plausible models of habituation for single-cell learning. Current Biology, 19 Nov 2024 (computational modelling). https://doi.org/10.1016/j.cub.2024.10.041 [Verified]
40. Gumuskaya, G., et al.; Levin, M. (Tufts) — Motile Living Biobots Self-Construct from Adult Human Somatic Progenitor Seed Cells. Advanced Science, 30 Nov 2023. https://doi.org/10.1002/advs.202303575 [Verified]
41. Adamala laboratory (University of Minnesota) — a bottom-up synthetic cell performing selection, genome replication, growth and genetically encoded division ('SpudCell'). bioRxiv preprint, July 2026 (DOI 10.64898/2026.07.01.735724); peer-review status unconfirmed; authors note consolidation and scaling limits. https://www.soci.org/news/2026/7/researchers-create-functioning-synthetic-cell [Secondary source]
42. Rosas, F. E., Mediano, P. A. M., et al. — Reconciling emergences: An information-theoretic approach to identify causal emergence in multivariate data. PLoS Computational Biology 16(12): e1008289, Dec 2020. https://doi.org/10.1371/journal.pcbi.1008289 [Verified]
43. Mediano, P. A. M., Rosas, F. E., Luppi, A. I., et al. — Greater than the parts: a review of the information decomposition approach to causal emergence. Philosophical Transactions of the Royal Society A 380(2227): 20210246, 2022. https://pmc.ncbi.nlm.nih.gov/articles/PMC9125226 [Verified]
44. Hoel, E. P., Albantakis, L., Tononi, G. — Quantifying causal emergence shows that macro can beat micro. PNAS 110(49), 2013. https://pmc.ncbi.nlm.nih.gov/articles/PMC3856819 [Partially verified]
45. Hoel, E. — Causal Emergence 2.0: Quantifying emergent complexity. arXiv 2503.13395 (March–April 2025), preprint. https://arxiv.org/abs/2503.13395 [Verified]
46. Eberhardt, F., Lee, L. L. — Causal Emergence: When Distortions in a Map Obscure the Territory. Philosophies 7(2): 30, 2022. https://doi.org/10.3390/philosophies7020030 [Verified]
47. Lau, E., Furman, Z., Wang, G., Murfet, D., Wei, S. — The Local Learning Coefficient: A Singularity-Aware Complexity Measure. AISTATS 2025, PMLR 258: 244–252; arXiv 2308.12108 (2023). Open estimator library: devinterp. https://proceedings.mlr.press/v258/lau25a.html [Verified]
48. Cullen, Estan-Ruiz, Danait, Li — A Basin-Selection Perspective on Grokking via Singular Learning Theory. arXiv 2603.01192 (March 2026; v3 May 2026), preprint. A source investigation cited an earlier title. https://arxiv.org/abs/2603.01192 [Verified]
49. Toker, D., et al. — Consciousness is supported by near-critical slow cortical electrodynamics. PNAS 119(7), 2022. https://doi.org/10.1073/pnas.2024455119 [Verified]
50. Maschke, C., et al.; Blain-Moraes, S. — Critical dynamics in spontaneous EEG predict anesthetic-induced loss of consciousness and perturbational complexity. Communications Biology 7, August 2024. https://doi.org/10.1038/s42003-024-06613-8 [Verified]
51. Touboul, J., Destexhe, A. — Power-law statistics and universal scaling in the absence of criticality. Physical Review E 95: 012413, 2017. https://doi.org/10.1103/PhysRevE.95.012413 [Verified]
52. Priesemann, V., et al. — Spike avalanches in vivo suggest a driven, slightly subcritical brain state. Frontiers in Systems Neuroscience 8: 108, 2014. https://doi.org/10.3389/fnsys.2014.00108 [Verified]
53. Beggs, J. M., Timme, N. — Being critical of criticality in the brain. Frontiers in Physiology 3: 163, 2012. https://www.frontiersin.org/articles/10.3389/fphys.2012.00163/pdf [Verified]
54. Clark, A., Chalmers, D. — The Extended Mind. Analysis 58(1): 7–19, 1998. https://web-archive.southampton.ac.uk/cogprints.org/1093/3/EXTENDED.pdf [Verified]
55. Hutchins, E. — Cognition in the Wild. MIT Press, 1995. https://mitpress.mit.edu/9780262581462/cognition-in-the-wild/ [Verified]
56. [Large collaboration; author list omitted] — Quantum error correction below the surface code threshold. Nature, online 9 December 2024 (105-qubit processor; logical error roughly halving with each increase in code distance). https://doi.org/10.1038/s41586-024-08449-y [Verified]
57. Durant, F., Morokuma, J., Fields, C., Williams, K., Adams, D. S., Levin, M. — Long-Term, Stochastic Editing of Regenerative Anatomy via Targeting Endogenous Bioelectric Gradients. Biophysical Journal 112(10): 2231–2243, 2017. https://doi.org/10.1016/j.bpj.2017.04.011 [Verified]

58. Levin, M. — Bioelectric networks: the cognitive glue enabling evolutionary scaling from physiology to mind. Animal Cognition 26, 2023. https://doi.org/10.1007/s10071-023-01780-3 [Verified]
59. Fruchart, M., Hanai, R., Littlewood, P. B., Vitelli, V. — Non-reciprocal phase transitions. Nature 592: 363–369, 15 April 2021. https://doi.org/10.1038/s41586-021-03375-9 [Verified]
60. Greilich, A., et al. (TU Dortmund, Ioffe Institute) — Robust continuous time crystal in an electron–nuclear spin system. Nature Physics 20: 631–636, 2024. https://doi.org/10.1038/s41567-023-02351-6 [Verified]
61. Zhao, H., Smalyukh, I. I. (University of Colorado Boulder) — Space-time crystals from particle-like topological solitons. Nature Materials, 4 September 2025. https://doi.org/10.1038/s41563-025-02344-1 [Verified]
62. Switzer, E. R., et al. — Realization of two-dimensional discrete time crystals with anisotropic Heisenberg coupling. Nature Communications 17: 605, 2026. https://doi.org/10.1038/s41467-025-67787-1 [Verified]
63. Sharma, A., Czégel, D., Lachmann, M., Kempes, C. P., Walker, S. I., Cronin, L. — Assembly theory explains and quantifies selection and evolution. Nature 622: 321–328, October 2023. https://doi.org/10.1038/s41586-023-06600-9 [Verified]
64. Uthamacumaran, A., Abrahão, F. S., Kiani, N. A., Zenil, H. — On the salient limitations of the methods of assembly theory and their classification of molecular biosignatures. npj Systems Biology and Applications 10: 82, 2024. See also Hazen et al., J. R. Soc. Interface 2024; Jaeger, J. Mol. Evol. 2024. https://doi.org/10.1038/s41540-024-00403-y [Verified]
65. Shumailov, I., et al. — AI models collapse when trained on recursively generated data. Nature, 24 July 2024; no measured real-world collapse reported as of 2026 (UTS Centre for Media Transition review, Sept 2026). https://www.cs.ox.ac.uk/news/2356-full.html [Verified]
66. Bestvater, S., Smith, A., TerBush, Baronavski, Chavda (Pew Research Center) — How Much of the Internet Is Written With AI? Pew Research Center, 20 August 2026. https://www.pewresearch.org/data-labs/2026/08/20/how-much-of-the-internet-is-written-with-ai/ [Verified]
67. Fiebelkorn, I. C., Pinsk, M. A., Kastner, S. — A Dynamic Interplay within the Frontoparietal Network Underlies Rhythmic Spatial Attention. Neuron 99(4): 842–853, 2018. https://pmc.ncbi.nlm.nih.gov/articles/PMC6474777 [Verified]
68. Landau, A. N., Fries, P. — Attention Samples Stimuli Rhythmically. Current Biology 22(11): 1000–1004, 2012. https://doi.org/10.1016/j.cub.2012.03.054 [Verified]
69. Markus, A., Shamay-Tsoory, S. G. — Hyperscanning: from inter-brain coupling to causality. Frontiers in Human Neuroscience 18: 1497034, November 2024 (perspective; evidence described as correlational). https://doi.org/10.3389/fnhum.2024.1497034 [Verified]
70. Burgess, A. P. — On the interpretation of synchronization in EEG hyperscanning studies: a cautionary note. Frontiers in Human Neuroscience 7: 881, 2013. https://pmc.ncbi.nlm.nih.gov/articles/PMC3870947 [Verified]
71. Tessler, M. H., Bakker, M. A., et al. — AI can help humans find common ground in democratic deliberation. Science 386, 18 October 2024 (N = 5,734 in the main study). https://doi.org/10.1126/science.adq2852 [Verified]
72. Du, Y., van Mastrigt, R., Veenstra, J., Coulais, C. — Metamaterials that learn to change shape. Nature Physics 22: 784–790, April 2026. https://doi.org/10.1038/s41567-026-03226-2 [Verified]
73. Melanson, D., et al. — Thermodynamic computing system for AI applications. Nature Communications 16, 2025 (eight all-to-all RLC cells performing Gaussian sampling and matrix inversion). https://doi.org/10.1038/s41467-025-59011-x [Verified]
74. Kagan, B. J., et al. (Cortical Labs) — In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. Neuron, 12 Oct 2022. The 'sentience' framing was contested by independent neuroscientists. https://doi.org/10.1016/j.neuron.2022.09.001 [Verified]
75. Tyler, W. J. — Ultrasonic Brain Computer Interfaces for Enhancing Human-Machine Cognition. arXiv 2604.00349, April 2026, preprint (review). Cited in place of a source investigation's 2026 ultrasound-interface reference that could not be located. https://arxiv.org/abs/2604.00349 [Verified]
76. Wright, L. G., Onodera, T., Stein, M. M., et al.; McMahon, P. L. (Cornell) — Deep physical neural networks trained with backpropagation. Nature 601: 549–555, January 2022. https://doi.org/10.1038/s41586-021-04223-6 [Verified]
77. Singh, M. P. — Information-Driven Interaction-Oriented Programming: BSPL, the Blindingly Simple Protocol Language. AAMAS 2011. https://www.csc2.ncsu.edu/faculty/mpsingh/papers/mas/AAMAS-11-IBIOP.pdf [Verified]
78. Zhan, Z. (Caltech) — Distributed Acoustic Sensing Turns Fiber-Optic Cables into Sensitive Seismic Antennas. Seismological Research Letters 91(1): 1–15, 2020. https://doi.org/10.1785/0220190112 [Verified]
79. Sipling, C., Zhang, Y.-H., Di Ventra, M. — A Critical Assessment of the Brain Criticality Hypothesis. arXiv 2604.21071 (v3, May 2026), preprint. https://arxiv.org/abs/2604.21071 [Verified]
80. Willis, Zhao, Du, Leibo — Evaluating Collective Behaviour of Hundreds of LLM Agents. arXiv 2602.16662 (Feb 2026), preprint. https://arxiv.org/abs/2602.16662 [Verified]

81. Zomer, N., De Domenico, M. — Unraveling the emergence of collective behavior in networks of cognitive agents. npj Artificial Intelligence 2: 36, 2026. https://doi.org/10.1038/s44387-026-00091-5 [Verified]
82. Koepke, Zverev, Ginosar, Efros et al. — Back into Plato's Cave: Examining Cross-modal Representational Convergence at Scale. arXiv 2604.18572 (2026), preprint. https://www.alphaxiv.org/abs/2604.18572 [Verified]
83. Cross-brain transfer of high-performance intracortical speech and handwriting BCIs. bioRxiv, 12 January 2026, preprint (five participants, 48.9 hours of data). https://www.biorxiv.org/content/10.64898/2026.01.12.699110v1 [Verified]
84. Hughes, E., Dennis, M., Parker-Holder, J., et al. — Position: Open-Endedness is Essential for Artificial Superhuman Intelligence. ICML 2024, PMLR 235; arXiv 2406.04268. https://arxiv.org/abs/2406.04268 [Verified]
85. W3C — Verifiable Credentials Data Model v2.0. W3C Recommendation, 15 May 2025. https://www.w3.org/TR/vc-data-model-2.0/ [Verified]
86. Kasselman, P., et al. — AI Identity Management System (AIMS), draft-ietf-wimse-aims-00. IETF WIMSE working group, Informational draft, 15 September 2026. Section 10.1 places translation of mission into authorization requirements out of scope. https://datatracker.ietf.org/doc/html/draft-ietf-wimse-aims [Verified]
87. Bansal, G., et al. — Magentic Marketplace: an open-source environment for studying agentic markets. arXiv 2510.25779, 27 Oct 2025. https://arxiv.org/abs/2510.25779 [Verified]
88. Adobe — Holiday 2025 online shopping recap: AI-referred traffic to US retail sites up 693% year on year. Adobe Analytics, 7 January 2026. https://news.adobe.com/news/2026/01/adobe-holiday-shopping-season [Verified]
89. Prindle, A., Liu, J., Asally, M., Ly, S., Garcia-Ojalvo, J., Süel, G. M. — Ion channels enable electrical communication in bacterial communities. Nature 527: 59–63, 2015. https://doi.org/10.1038/nature15709 [Verified]
90. Yang, C.-Y., Bialecka-Fornal, M., et al.; Süel, G. M. (UC San Diego) — Encoding membrane-potential-based memory within a microbial community. Cell Systems, 27 April 2020. https://today.ucsd.edu/story/they-remember-communities-of-microbes-found-to-have-working-memory [Secondary source]
91. Coles, P. J., Szczepanski, C., Melanson, D., Donatella, K., Martinez, A. J., Sbahi, F. — Thermodynamic AI and the fluctuation frontier. arXiv 2302.06584 (v3, June 2023), preprint. https://arxiv.org/abs/2302.06584 [Verified]
92. Normal Computing — Tape-out of CN101 thermodynamic computing chip. 12 August 2025 (company announcement). https://normalcomputing.com/blog/normal-computing-announces-tape-out-of-worlds-first-thermodynamic-computing-chip [Secondary source]
93. Cai, H., et al. (Guo, F., Indiana University) — Brain organoid reservoir computing for artificial intelligence. Nature Electronics, December 2023. https://doi.org/10.1038/s41928-023-01069-w [Verified]
94. Liu, T., Philamore, H., Ward-Cherrier, B. (University of Bristol) — Encoding Tactile Stimuli for Organoid Intelligence in Braille Recognition (FinalSpark Neuroplatform). arXiv 2508.20850 (2025), preprint. https://arxiv.org/abs/2508.20850 [Verified]
95. Sharma, A., Ng, T.-K. M., Parrilla Gutierrez, J. M., Yibin, J., Cronin, L. — A programmable hybrid digital-chemical information processor based on the Belousov–Zhabotinsky reaction. Nature Communications 15, 2024. https://doi.org/10.1038/s41467-024-45896-7 [Verified]
96. Lee, R., et al.; Chou, L. Y. T. (University of Toronto) — Reusable enzyme-driven DNA building blocks for molecular machine learning. Journal of the American Chemical Society, September 2024 (DOI 10.1021/jacs.4c07221). A source investigation's citation to a different journal was incorrect. https://doi.org/10.1021/jacs.4c07221 [Partially verified]
97. Morrell, M. C., Elliott, A., Grier, D. G. (New York University) — Nonreciprocal wave-mediated interactions power a classical time crystal. Physical Review Letters, DOI 10.1103/zjzk-t81n (2025–2026); arXiv 2504.15495. Acoustically levitated beads sustaining oscillation for hours without periodic drive. https://arxiv.org/abs/2504.15495 [Verified]
98. O'Connor, A., Montenegro, V., Albarelli, F., Paris, M. G. A., Bayat, A., Genoni, M. G. — Attaining Quantum Sensing Enhancement from Monitored Dissipative Time Crystals. arXiv 2508.15448 (August 2025; v2 January 2026), preprint; theoretical. https://arxiv.org/abs/2508.15448 [Verified]
99. TechCrunch — report on a withdrawn claim that a language model had solved ten open Erdős problems; the problems' curator described the claim as 'a dramatic misinterpretation' after the 'solutions' proved to be literature retrieval. 19 October 2025 (news report; cited by outlet and date). [Secondary source]
100. EPFL — News item: Viazovska's sphere-packing proofs in dimensions 8 and 24 formalized in Lean by a commercial automated formalization system [developer and system name omitted]. EPFL news, 11 March 2026. https://actu.epfl.ch/news/prof-viazovska-s-proofs-of-sphere-packing-formaliz/ [Verified]
101. Unite.AI — Report: a Lean 4 proof of the bounded prime gaps (246) theorem produced by an automated prover from a human blueprint [developer and system name omitted]. August 2026 (secondary report). https://unite.ai/axiom-maths-ai-verifies-the-246-prime-gaps-theorem-in-lean/ [Secondary source]

102. TechCrunch — report that two experimental reasoning systems reached gold-medal standard at the 2025 International Mathematical Olympiad (35/42), one graded and certified by IMO coordinators. 21 July 2025 (news report; cited by outlet and date). [Secondary source]
103. Zhang, J., Hu, S., Lu, C., Lange, R., Clune, J. — Darwin Gödel Machine: Open-Ended Evolution of Self-Improving Agents. arXiv 2505.22954, May 2025; published at ICLR 2026 after peer review (v3, March 2026). Self-modifying coding agent with an archive of variants; documented safety observations. https://arxiv.org/abs/2505.22954 [Verified]
104. Lu, C., Lange, R., Yamada, Y., Hu, S., Foerster, J., Ha, D., Clune, J. — Towards end-to-end automation of AI research. Nature 651: 914–919, 25 March 2026. https://ora.ox.ac.uk/objects/uuid:92d16200-fbe1-4564-b4d4-694e96223b39 [Verified]
105. Guess, A. M., et al. — How do social media feed algorithms affect attitudes and behavior in an election campaign?; Nyhan, B., et al. — Like-minded sources on Facebook are prevalent but not polarizing. Science and Nature, 27 July 2023 (US 2020 Facebook and Instagram Election Studies). https://doi.org/10.1038/s41586-023-06297-w [Verified]
106. Bail, C. A., et al. — Exposure to opposing views on social media can increase political polarization. PNAS 115(37): 9216–9221, 2018. https://doi.org/10.1073/pnas.1804840115 [Verified]
107. Salvi, F., Horta Ribeiro, M., Gallotti, R., West, R. — On the conversational persuasiveness of [a large language model; model name abridged from title]. Nature Human Behaviour, 2025 (preregistered; 900 participants; personalization raised odds of opinion shift by 81% relative to human–human debate). https://doi.org/10.1038/s41562-025-02194-6 [Verified]
108. Hackenburg, K., Tappin, B. M., Hewitt, L., et al. — The levers of political persuasion with conversational artificial intelligence. Science 390(6777), 4 Dec 2025; arXiv 2507.13919. https://arxiv.org/abs/2507.13919 [Verified]
109. Lin, H., Czarnek, G., Lewis, S., et al. (Rand, D. G.) — Persuading voters using human–artificial intelligence dialogues. Nature, 4 December 2025. https://doi.org/10.1038/s41586-025-09771-9 [Verified]
110. Kirk, H. R., Vidgen, B., Röttger, P., Hale, S. A. — The benefits, risks and bounds of personalizing the alignment of large language models to individuals. Nature Machine Intelligence 6(4): 383–392, April 2024. https://doi.org/10.1038/s42256-024-00820-y [Verified]
111. Costello, T. H., Pennycook, G., Rand, D. G. — Durably reducing conspiracy beliefs through dialogues with AI. Science, 13 September 2024. https://doi.org/10.1126/science.adq1814 [Verified]
112. Kunz, E. M., Abramovich Krasa, B., Willett, F. R., et al. (Stanford) — Inner speech in motor cortex and implications for speech neuroprostheses. Cell, online 14 August 2025. https://www.ninds.nih.gov/news-events/news/press-releases/decoding-inner-speech-brain-signals [Verified]
113. Kaifosh, P., Reardon, T. R., et al. — A generic non-invasive neuromotor interface for human-computer interaction. Nature 645(8081): 702–711, September 2025 (online July 2025). https://doi.org/10.1038/s41586-025-09255-w [Verified]
114. Card, N. S., et al. — An Accurate and Rapidly Calibrating Speech Neuroprosthesis. New England Journal of Medicine 391(7): 609–618, 15 Aug 2024. https://doi.org/10.1056/NEJMoa2314132 [Verified]
115. Card, N. S., Singer-Clark, T., Peracha, H., et al. — Long-term independent use of an intracortical brain–computer interface for speech and cursor control. Nature Medicine, 15 June 2026. https://www.natureasia.com/en/info/press-releases/detail/9358 [Secondary source]
116. Salazar-Gomez, A. F., DelPreto, J., Gil, S., Guenther, F. H., Rus, D. — Correcting robot mistakes in real time using EEG signals. IEEE ICRA 2017. https://doi.org/10.1109/ICRA.2017.7989777 [Verified]
117. Lévy, J., Rapin, J., Banville, H., et al. — Brain-to-Text Decoding: A Non-invasive Approach via Typing. arXiv 2502.17480, February 2025 (35 participants; character error rate 32% MEG, 67% EEG; offline). https://arxiv.org/abs/2502.17480 [Verified]
118. Tang, J., LeBel, A., Jain, S., Huth, A. G. — Semantic reconstruction of continuous language from non-invasive brain recordings. Nature Neuroscience, 1 May 2023. https://doi.org/10.1038/s41593-023-01304-9 [Partially verified]
119. Horikawa, T. (NTT Communication Science Laboratories) — Mind captioning: generating descriptive text from brain activity. Science Advances, 5 November 2025. https://www.scientificamerican.com/article/ai-decodes-visual-brain-activity-and-writes-captions-for-it/ [Secondary source]
120. Jiang, L., Stocco, A., Losey, D. M., Abernethy, J. A., Prat, C. S., Rao, R. P. N. — BrainNet: A Multi-Person Brain-to-Brain Interface for Direct Collaboration Between Brains. Scientific Reports 9: 6115, 2019. https://doi.org/10.1038/s41598-019-41895-7 [Verified]
121. Dikker, S., et al. — Brain-to-Brain Synchrony Tracks Real-World Dynamic Group Interactions in the Classroom. Current Biology 27(9): 1375–1380, 2017. https://doi.org/10.1016/j.cub.2017.04.002 [Verified]
122. Réveillé, C., Vergotte, G., Perrey, S., Bosselut, G. — Using interbrain synchrony to study teamwork: A systematic review and meta-analysis. Neuroscience & Biobehavioral Reviews 159: 105593, April 2024.

https://doi.org/10.1016/j.neubiorev.2024.105593 [Verified]

123. Brosler, S. C., … Chang, E. F. (UCSF) — Simultaneous speech and gesture decoding for multimodal communication in paralysis. Nature Neuroscience, 14 September 2026. https://www.nih.gov/news-events/news-releases/neuroprosthesis-paralysis-enables-simultaneous-speech-body-language [Secondary source]
124. Cogitate Consortium (Melloni, L., Mudrik, L., Pitts, M., et al.) — Adversarial testing of global neuronal workspace and integrated information theories of consciousness. Nature, 2025 (256 participants; seven laboratories; neither theory vindicated). https://doi.org/10.1038/s41586-025-08888-1 [Verified]
125. Marra, G., et al. (National Physical Laboratory) — Optical interferometry-based array of seafloor environmental sensors using a trans-oceanic submarine cable. Science 376: 874–879, 19 May 2022. https://www.research.ed.ac.uk/en/publications/optical-interferometry-based-array-of-seafloor-environmental-sens/ [Verified]
126. Allen, R. M., Barski, M., Berman, D., et al. — Global earthquake detection and warning using smartphones [title paraphrased to omit a product name]. Science, 17 July 2025, pp. 254–259; coverage expansion to more than 200 countries and regions announced by the operator, 15 October 2026. https://doi.org/10.1126/science.ads4779 [Verified]
127. Hutchison, C. A., et al. — Design and synthesis of a minimal bacterial genome. Science 351(6280): aad6253, 2016 (JCVI-syn3.0). https://www.jcvi.org/publications/design-and-synthesis-minimal-bacterial-genome [Verified]
128. Pelletier, J. F., et al. — Genetic requirements for cell division in a genomically minimal cell. Cell 184(9): 2430–2440, 2021 (JCVI-syn3A). https://www.jcvi.org/publications/genetic-requirements-cell-division-genomically-minimal-cell [Verified]
129. Huang, H.-Y., Broughton, M., Cotler, J., et al. — Quantum advantage in learning from experiments. Science 376: 1182–1186, 2022. https://doi.org/10.1126/science.abn7293 [Verified]
130. Qu, D., Lin, Q., Xiao, L., Zhan, X., Xue, P. — Experimental violation of a Bell-like causal inequality in a photonic quantum switch. Science Advances 12(32): eaee9271, 7 August 2026 (received December 2025, accepted June 2026); preprint arXiv 2508.04643, August 2025, under the title 'Experimental device-independent certification of indefinite causal order'. Inequality value 1.8090 ± 0.0024 against the definite-order bound 1.75 (24σ); locality and freedom-of-choice loopholes addressed; fair-sampling assumption retained (global detection efficiency ≈12%). https://doi.org/10.1126/sciadv.aee9271 [Verified]
131. Richter, C. M. D., Antesberger, M., Cao, H., Walther, P., Rozema, L. A. — Toward an experimental device-independent verification of indefinite causal order. PRX Quantum 7: 010354, 2026; preprint arXiv 2506.16949, June 2025. Independent photonic quantum-switch violation of the same Bell-like causal inequality: 1.8328 ± 0.0045 against the bound 1.75 ($\approx 18\sigma$); locality loophole open. https://doi.org/10.1103/5t2y-ddmt [Verified]
132. VanRullen, R. — Perceptual Cycles. Trends in Cognitive Sciences 20(10): 723–735, 2016. https://doi.org/10.1016/j.tics.2016.07.006 [Verified]
133. Babcock, N. S., et al. (Kurian, P., Howard University; Chergui, M., EPFL) — Ultraviolet Superradiance from Mega-Networks of Tryptophan in Biological Architectures. Journal of Physical Chemistry B, 2024. https://doi.org/10.1021/acs.jpcb.3c07936 [Verified]
134. Patwa, H., Babcock, N. S., Kurian, P. — Quantum-enhanced photoprotection in neuroprotein architectures emerges from collective light–matter interactions. Frontiers in Physics 12: 1387271, August 2024 (computational models). A source investigation cited this paper under an incorrect title. https://doi.org/10.3389/fphy.2024.1387271 [Verified]
135. Kurian, P. — Room-temperature superradiance from quantum optical networks in protein fibers (invited talk; superradiant lifetimes of no more than a few picoseconds). APS March Meeting 2026, abstract. https://meetings-archive.aps.org/smt/2026/mar-y19/1/ [Verified]
136. Khan, S., et al.; Wiest, M. C. (Wellesley College) — Microtubule-Stabilizer Epothilone B Delays Anesthetic-Induced Unconsciousness in Rats. eNeuro 11(8), 15 Aug 2024. https://doi.org/10.1523/ENEURO.0291-24.2024 [Verified]
137. Tegmark, M. — Importance of quantum decoherence in brain processes. Physical Review E 61: 4194, 2000. https://doi.org/10.1103/PhysRevE.61.4194 [Verified]
138. Shwartz-Ziv, R., Tishby, N. — Opening the Black Box of Deep Neural Networks via Information. arXiv 1703.00810, 2017. https://arxiv.org/abs/1703.00810 [Verified]
139. Saxe, A. M., et al. — On the information bottleneck theory of deep learning. ICLR 2018; Journal of Statistical Mechanics 2019. https://research.ibm.com/publications/on-the-information-bottleneck-theory-of-deep-learning [Verified]
140. Gündüz, D., Qin, Z., Estella Aguerri, I., et al. — Beyond Transmitting Bits: Context, Semantics, and Task-Oriented Communications. IEEE Journal on Selected Areas in Communications 41(1): 5–41, Jan 2023. https://arxiv.org/abs/2207.09353 [Verified]

APPENDIX B

Methodological notes

B.1 Classification

Classes follow the founding mandate and were used by all four investigations. **Class I — Emerging reality:** a phenomenon supported by credible, peer-reviewed observations but not yet fully understood or unified. **Class II — Theoretical frontier:** a hypothesis grounded in known science but not yet experimentally demonstrated. **Class III — Radical conjecture:** a speculative concept that would require substantial revision of existing assumptions. **Class IV — Conceptual invention:** a new explanatory framework or category proposed by an EONORYX investigation. Classes are not rankings of importance or truth. An entry may carry two classes when a Class I observation is paired with a Class IV framework proposed to explain it. Class I requires credible peer-reviewed publication; replication is recorded by the evidence grade, not the class.

B.2 Evidence grades

E3 replicated or independently confirmed; **E2** bounded observation — peer-reviewed, or a preprint so labelled — not yet independently replicated or generalized; **E1** precursors only — related findings exist, the specific claim is untested; **E0** no evidence, a proposal. Grades attach to the entry's specific claim, not to its field: an entry that proposes a new quantity whose ingredients are measured but whose own value has never been measured is graded E1, however strong the surrounding field (DR-22, DR-26, DR-27). An entry that extends an observed phenomenon to a new setting carries the grade of the bounded observation, E2, and its novelty field states the extension that remains untested (DR-17, DR-24, DR-25, DR-34). Preprints are graded no higher than E2 and are labelled: an observation available only as a preprint carries the grade label "preprint, peer review pending", and a Class I entry that rests on a preprint alone is additionally marked provisional (DR-30). Grades are re-assessed monthly for monitor entries and on any flagged result; every change is logged with its reason.

B.3 Status

Priority — enters a numbered programme with a preregistration and a day-90 decision. **Monitor** — tracked by the Horizon Monitor with a defined trigger for promotion. **Hold** — awaits a partner laboratory or a proponent's specification of the parameters needed to test it. **Archived** — rejected or superseded; retained with the reason and the evidence that would reopen it.

B.4 The precedent audit

Every coined term is tested against one question: does it pick out a measurable claim that its nearest precedent does not? Outcomes are *withdrawn* (the precedent's name is used), *alias* (recorded for provenance; the entry carries a plainer name) or *retained* (the added claim is stated in the entry's novelty field). The audit is repeated whenever a new precedent is found.

B.5 Scoring

This charter does not carry forward numerical novelty, plausibility or testability scores from the source investigations, because none stated how its scores were derived. The qualitative assessments in each registry entry (novelty assessment; competing explanations) replace them. A future edition may introduce a scored rubric once inter-rater reliability has been measured on the register.

B.6 Verification tags

Each reference carries one of: *Verified* — retrieved and checked against the primary record or an institutional copy; *Partially verified* — existence and core claim confirmed, some bibliographic detail unconfirmed; *Secondary source* — the primary record could not be reached and a reputable secondary report is cited; *Not independently verified* — cited by a source investigation, retained for provenance only, not relied upon; *Source investigation* — one of the four commissioned drafts. Verification was performed in October–November 2026; a reference's tag states what was checked, not that the work is correct.

B.7 Preregistration template

Each programme's preregistration states: the registry entries tested; the specific hypotheses; the measures and their estimators; the controls and surrogates; sample sizes and seeds; the analysis plan; the decision rule (continue, reformulate, close) with numerical thresholds; the designated critic; and the data and code release plan. Amendments after pilots are appended, dated and explained.

APPENDIX C

Definitions

Capacity
A function a system performs, defined independently of the system's self-description: a learned task, a verified result, a shared description, an institutional commitment.

Carrier
The physical or computational substrate that currently holds a capacity: an agent, a population, a material, a tissue, an operator, a protocol.

Carrier question
What remains invariant, and how it is measured, when the carrier of a capacity changes.

Coherence Index
Mean pairwise mutual information between individuals' descriptions of a common set of verifiable items, conditional on the degree of personalization of their information environment.

Commons Decoherence
A regime in which individually adaptive interpreters improve individual accuracy while degrading a population's capacity to coordinate on a common description, at equal evidence and aligned objectives.

Competence migration
A measured redistribution of the task-relevant information and mechanisms from identifiable participants into a persistent interaction structure or environment.

Discovery Registry
EONORYX's versioned, machine-readable register of candidate phenomena with classification, evidence grade, provenance, falsifiers, experiments and status.

Displacement
The hypothesized movement of a capacity from its historical carrier to a new one; the Unnamed Transition comprises four.

Dossier
A preregistered EONORYX research report with data, code, the designated critic's report and a stated decision.

Emergence
The dimension asking under what conditions new capabilities, properties or forms of organization arise that are not present in, or predictable from, the components.

Evidence grade
E0–E3; see Appendix B.

Family
A group of registry entries sharing an experimental grammar — a single design that could in principle test all of them.

Investigations A–D
The four commissioned research drafts synthesized in this charter; see Appendix D.

Kill criterion
A result, specified before data are collected, that closes a programme.

Latent Concordance
Calibrated agreement in the relational structure of representations between independently trained learners, hypothesized to be controlled by the structure of the world they model.

Local learning coefficient
The singularity-aware complexity measure of singular learning theory; a candidate order parameter for collective learning when estimated over a coupled ensemble.

Order parameter
A measurable quantity whose sharp change marks a transition between regimes.

Persistence
The dimension asking how information, competence, memory, identity or function survives the replacement of its carrier.

Precedent audit
The procedure that tests every coined term against its nearest existing concept; see Appendix B.

Recovery ratio (R)

$(L_d - L_r)/(L_d - L_0)$: the fraction of damage-induced loss removed by an unlabelled recovery period.

Regenerative task memory

Recovery of a learned relation in a physical learning network after damage, without teacher replay or archived parameters.

Synergy

In partial information decomposition, the information about a target carried only by the joint state of several sources.

Transformation

The dimension asking when a system crosses a threshold beyond which existing categories stop predicting its behaviour.

Unnamed Transition

The Class IV hypothesis that four capacities — cognition, knowledge, reality and computation — are simultaneously loosening from their historical carriers.

Verification

The fourth dimension, asking under what conditions a knowledge-producing system retains the ability to detect that it is wrong.

Verification Gap

Newly validated claims minus newly witnessed claims in a field over a period.

Witness; Witness Index

A claim is witnessed if a human can reconstruct its justification to the field's standard; the index is the witnessed fraction of validated claims.

APPENDIX D

Provenance

D.1 Commission

Master Research Mandate 001 (October–November 2026) commissioned independent investigations into phenomena that humanity has not yet learned to name, specifying five conceptual provocations, a classification scheme, a requirement of at least fifteen candidate phenomena and five priority hypotheses, and a required memorandum structure. Final Synthesis Mandate 002 (November 2026) commissioned this charter as a new synthesis of the four resulting investigations, with the additional requirements of a unified research architecture, a Discovery Registry, a staged research programme, a scientific-integrity chapter, a digital platform architecture and a founder's closing thesis. Both mandates were written and issued by the founder.

D.2 The four investigations

CODE	FORM RECEIVED	SCOPE	HOW IT WAS USED
A	Research memorandum (47 pages); structured candidate register (15 entries); source register (30 entries); preregistration template; two executed, seeded negative-control simulations with code and results; checksums	Functional continuity under replacement; five priority programmes; three items belonging to a private workstream, excluded from this charter	Register and protocols adopted in large part; evidence grades adopted; negative controls preserved in the archive and described in Chapter 9; coined terms audited
B	Research memorandum (24 pages) with 26 references	Physical learning, bioelectric morphogenesis, singular-learning-theory order parameter, non-reciprocal active matter, post-platform self-organization; three counter-theses	Mechanisms and intersections adopted; eight coined terms withdrawn or demoted after audit; several citations corrected or marked unverified; numerical scores not carried forward
C	Abbreviated reconstruction only (the full original text was not available to the synthesis)	Criticality proximity, temporal multiplexing, synthetic cell autonomy, intention and error decoding, capacity-limited information flow; ten further named candidates	Entries provisionally graded; specific claims independently checked where possible (one project confirmed as a 2026 preprint; one frequency claim unsupported; one analogy archived); to be revisited when the full text is supplied
D	Research memorandum (50 pages) with 103 verified references	Cognitive delocalization; latent concordance, synergy transition, Verification Gap, consensus decoherence, closed-loop interneural integration; eighteen candidates; a discovery laboratory; a 90-day programme	Architecture of displacements adopted and extended to four dimensions; verification practice adopted; one coined term withdrawn, several demoted; the "intelligence field" and "consciousness coupling" rejections shared with A

D.3 What the founder supplied, and what the systems did

The founder supplied the mandates, the selection of investigations, the decision to synthesize, the visual identity of the institution, the editorial direction (including the exclusion of private workstreams from the public document and the requirement that the publication carry one institutional identity), and final review. The assisted systems performed literature search and retrieval, drafting, verification of citations against primary records, reconciliation of terminology, composition of the text, generation of figures and tables from data, and the build of the publication. No assisted system performed a laboratory experiment; no human performed one either. The two executed computations are simulations of elementary probability, described in Chapter 9.

D.4 The internal research archive

The archive preserves, unmodified, the four source investigations as received; the mandates; the verification logs; the brand-extraction record; the editable master of this charter; and a provenance map stating which assisted system produced which investigation. The map is held internally so that provenance is recoverable; the public charter does not name systems, because its subject is the phenomena.

APPENDIX E

Version history

VERSION	DATE	STATUS	CHANGES
1.0	2 November 2026	Foundational edition, for review	First publication. Four-dimension architecture (Emergence, Persistence, Transformation, Verification). Discovery Registry v1.0: 35 active entries in nine families and an unaffiliated group; 9 archived. Eight priority programmes. 90-day programme and 2027–2036 horizon. Precedent audit of 24 coined terms. Methodology and assistance disclosure. Known limitations: Investigation C represented by an abbreviated reconstruction; no hypothesis yet tested; several source citations marked unverified (see the accompanying editorial note).
1.0.1	2 November 2026	Corrective release	Scientific-consistency correction before public release; no change to structure, programmes, authorship or design. DR-32: source confirmed as peer-reviewed (Science Advances, Aug 2026) and independently confirmed (PRX Quantum, 2026); grade E2 → E3, provisional marker removed, second reference added. DR-30: grade E1 → E2 with the preprint label, since the preprint reports the defined behaviour; provisional marker retained. DR-25: precedent and reference updated to the peer-reviewed 2026 publication. Evidence-grade E2 wording aligned with the preprint rule in B.2 (front matter, B.2, registry schema); B.2 extension rule now lists DR-34. References: 140.
1.1	planned, ~February 2027	After the 90-day programme	Re-graded entries; eight programme decisions; first Witness Index; DR-28 protocol; Investigation C revisited if its full text is supplied.
2.0	planned, ~November 2027	First annual revision	First-year results; external review record; first-year plan outcome.

Corrections to this edition may be submitted through the research platform described in Chapter 14. Every accepted correction will be recorded here with its reason.

EONORYX

READ THE HORIZON · EONORYX.COM