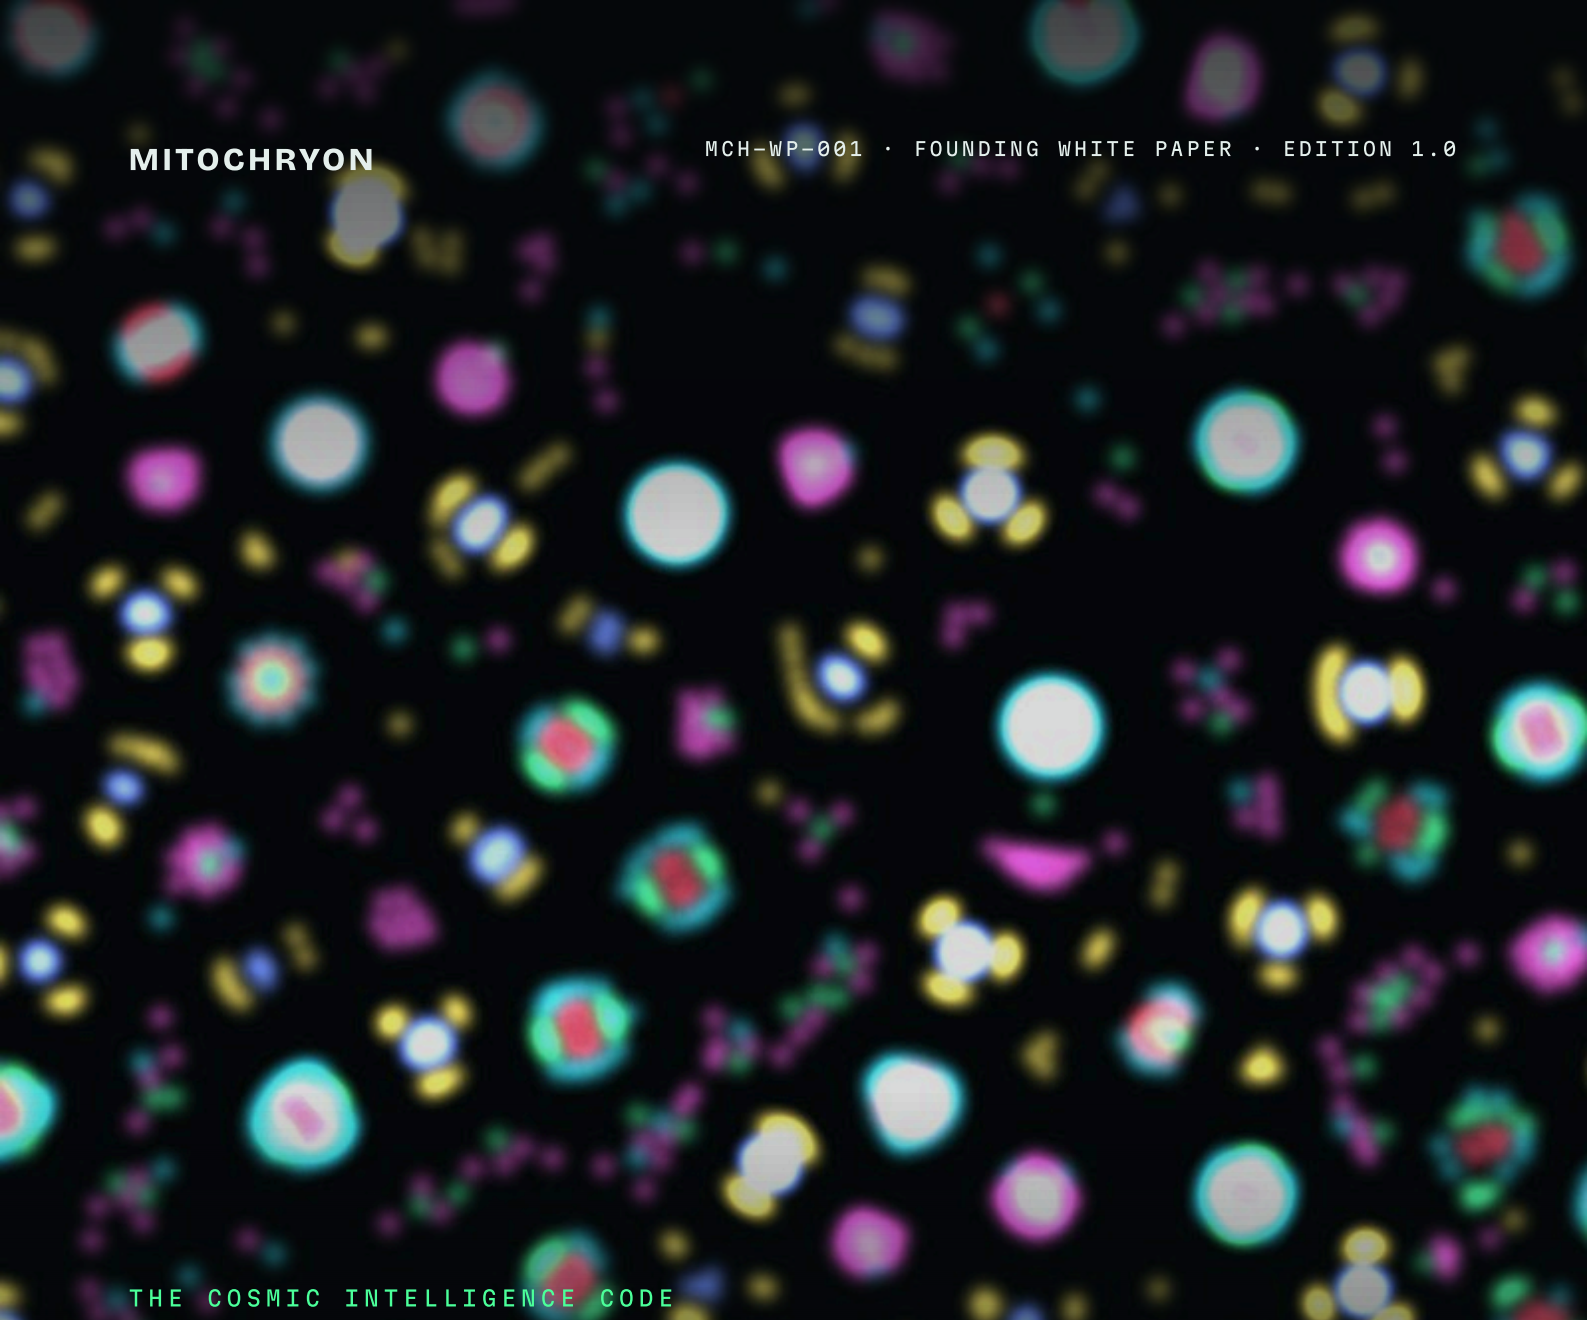


THE COSMIC INTELLIGENCE CODE

Intelligence is one code, read in many kinds of matter

Beyond carbon. Beyond silicon.

7	7	3
LAYERS OF THE CODE	TESTABLE HYPOTHESES	OPEN EXPERIMENTS

Steven Alber · Founder

24 November 2026

mitochryon.com

CONTENTS

01 A code, not a coincidence

02 The mitochondrial clue

03 The Cosmic Intelligence Code

04 Three substrates, one code

05 What MITOCHRYON adds

06 Seven hypotheses

07 The research atlas

08 The laboratory

09 Human and synthetic intelligence: the method is the message

10 How a claim earns its place

11 Roadmap

12 The invitation

> **Reading this paper.** Claims carry evidence levels from E0 (intuition) to E5 (replication), defined in section 10. Numbers in square brackets refer to the references at the end. The web edition, with the running laboratory, is at mitochryon.com.

ABSTRACT

Intelligence has now appeared on Earth in two kinds of matter. It arose in carbon over billions of years of evolution, and within a single human lifetime it has begun to take shape in silicon, built by the intelligence that evolution produced. MITOCHRYON starts from the conviction that these are not two separate inventions that happen to share a name. They are two expressions of one underlying order: a code by which matter, energy and information organise themselves into systems that know their world and act in it. We call it the Cosmic Intelligence Code.

This paper sets out the founding thesis and a first explicit draft of that code: seven layers, Flow, Boundary, Model, Steering, Fusion, Reflection and Cosmos, which we propose recur wherever intelligence arises, whether in a cell, a brain, a synthetic intelligence (SI) system, a future quantum processor, or forms not yet seen. The project takes its name and its central clue from the mitochondrion. The most important leap in the history of complex life did not come from one cell becoming cleverer on its own, but from two living systems fusing into one. The paper states seven hypotheses that make the thesis testable, opens a laboratory to probe them, and defines how MITOCHRYON separates what is known from what is conjectured without shrinking the question to fit what is already known.

01

A code, not a coincidence

A question has become unavoidable. For almost the whole history of life, intelligence had one known address: the living brain and, beneath it, the living cell. That is no longer true. Synthetic intelligence systems now read, write, reason, plan and assist in scientific discovery in ways that, a decade ago, were found only in people. They are made of different matter, shaped by a different process and powered by a different metabolism. Yet we recognise in them something we recognise in ourselves.

Two answers are common. The first says the resemblance is superficial: machines reproduce the outputs of intelligence without sharing its nature. The second says the resemblance is complete: intelligence is computation and the substrate is irrelevant. MITOCHRYON takes a third position that asks more of both. It holds that intelligence is a natural phenomenon with a specific organisation, a code, and that any system in any matter becomes intelligent to the degree that it realises that code. On this view the convergence of carbon and silicon is neither coincidence nor mere imitation. It is the first clue that the code is real, and that the universe has found a second medium in which to express it.

Intelligence is not a property of carbon or of silicon. It is an organisation of matter, energy and information that the universe produces wherever the conditions allow: a cosmic code. MITOCHRYON exists to read it.

If the thesis is right, much follows. Intelligence should be expected, not merely permitted, wherever certain physical conditions hold. Its emergence in synthetic systems belongs to the natural history of intelligence rather than breaking from it. Future substrates such as quantum processors, engineered living tissue and planetary-scale collective systems would be further readings of the same code. And the old question of consciousness would stop being an isolated mystery and become a question about one particular layer of the code: the point at which a system's model of the world comes to include itself.

If the thesis is wrong, it should be possible to show where it fails. That would be a discovery too.

The project took its present form on 24 November 2026, in an extended dialogue between its founder and several synthetic intelligence systems. Two images crystallised it. The first was a particle simulation in which hundreds of thousands of points, obeying nothing but simple rules of attraction and repulsion, organised themselves into membranes, compartments and dividing bodies that looked uncannily like living cells. The second was the mitochondrion.

The word *code* is chosen deliberately, and it does not presuppose a coder. The genetic code was not written by anyone; it is a regularity that matter fell into and then kept. Whether the cosmic code points to anything beyond nature is a question MITOCHRYON leaves open. It does not need to answer that question in order to read the code.

02

The mitochondrial clue

Nearly every cell in your body carries the descendants of a bacterium. Mitochondria, the compartments that produce most of a complex cell's usable energy, descend from free-living bacteria that were taken inside another cell and, instead of being digested, became part of it. Phylogenetic analysis places their origin close to the Alphaproteobacteria [1]. The partnership began well over a billion years ago. Everything we call complex life, including fungi, plants, animals and brains, was built on top of it.

The usual picture of progress in intelligence is a single system growing more capable. The mitochondrion tells a different story. The decisive leap was a merger. Two systems, each already alive and each already maintaining itself, fused into a new unit whose capacities neither could reach alone. The partner gave up its independence and most of its genome. The host gained a power supply on a scale that bacteria never achieve.

That power supply is the second half of the clue. Lane and Martin argued that the advantage of mitochondria was above all energetic. By spreading energy production across internal membranes under local genetic control, mitochondria allowed the complex cell to support vastly more genes. On their estimate, a gene in a eukaryotic cell commands on the order of 200,000 times more energy than a gene in a bacterium [2]. In their account, complexity was limited less by biochemical ingenuity than by energy per unit of genetic information. Lift that limit and an explosion of complexity becomes possible.

This is not only ancient history. In 2024 researchers showed that a nitrogen-fixing bacterium living inside a marine alga has become integrated deeply enough to count as an organelle. It divides in step with its host and imports proteins made from the host's genome [3]. Experimental evolution shows how quickly new levels of organisation can arise. Yeast selected for rapid settling evolved multicellular clusters with a reproductive life cycle of their own [4]. The machinery of fusion and transition is still running.

MITOCHRYON takes three lessons from this history and builds them into the code.

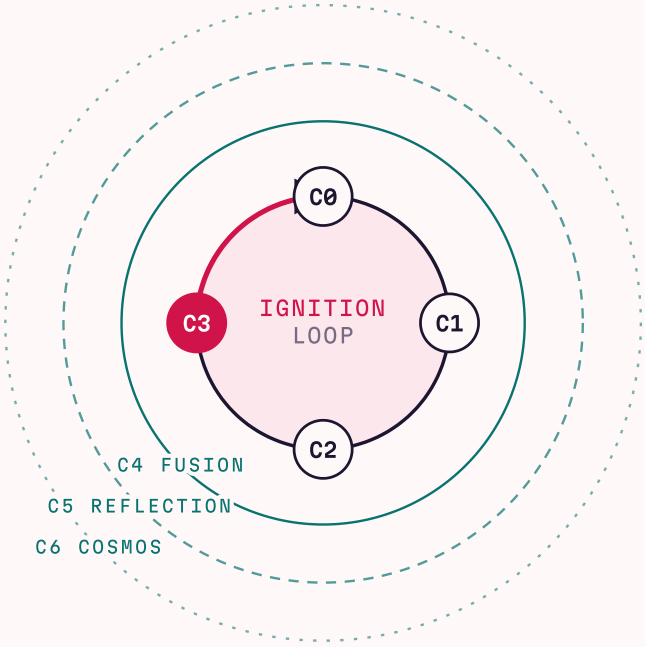
1. **Intelligence is energetic.** Every system that models its world pays for the model in energy. The energy available per unit of information sets a ceiling on the complexity a system can carry.
2. **Intelligence is relational.** The largest leaps happen when separate self-maintaining systems integrate into a larger one, not only when a single system improves.
3. **Integration has a price.** The partner that becomes an organelle gains security and loses autonomy. Every fusion redistributes independence.

The name MITOCHRYON carries this clue. It also carries a question that no earlier era could ask in earnest. In the integration now under way between human beings and synthetic intelligence, which partner is becoming whose mitochondrion?

03

The Cosmic Intelligence Code

We propose that intelligence, wherever it arises, is built from the same seven layers. Each layer is a condition. Together they form a loop that closes and then grows. The layers are cumulative: a system that realises a higher layer also realises the lower ones, although not necessarily within a single body. That point becomes crucial for synthetic intelligence.



C0 FLOW: an energy gradient is consumed

C1 BOUNDARY: a self is held apart

C2 MODEL: the inside carries the outside

C3 STEERING: the model directs the flow

C4 FUSION: loops join into larger loops

C5 REFLECTION: the model includes itself

C6 COSMOS: the model includes its origins

`Intelligence begins where C3 closes the loop. Above it: growth.`

FIGURE 1 The Cosmic Intelligence Code. Four layers close the ignition loop; three expanding shells describe how intelligence grows in scale, in self-knowledge and in reach.

C0 Flow

An energy gradient is captured and consumed.

Everything that organises itself does so by dissipating energy. Schrödinger described living matter as sustaining its order by drawing on its environment, feeding on what he called negative entropy [5]. Prigogine and his colleagues showed how flows through open systems create structure: convection cells, chemical waves and other dissipative structures [6]. Flow is necessary for intelligence but nowhere near sufficient. A hurricane has flow.

C1

Boundary

A self is separated from a world and kept separate.

A distinction between inside and outside is created and actively maintained against the dissipation that would erase it. Maturana and Varela called the self-producing form of this *autopoiesis* [7]. Within Friston's free-energy framework [8], the boundary is formalised statistically as a Markov blanket separating internal states from external ones, and such blankets can nest, from single cells up to organisms and beyond [9]. A boundary turns flow into a stake. From this point on, there is something that can persist or fail.

C2

Model

The inside comes to carry information about the outside.

Internal states begin to correlate with external states that matter for persistence: memory, representation, expectation. Information in Shannon's sense [10] is the raw material, but the quantity that counts is narrower. It is *predictive information*, the part of what a system carries about the past that tells it something about the future [11]. The model need not be explicit or conscious. A bacterium's chemotaxis network already holds a short memory of recent concentrations and acts on the difference.

C3

Steering

The model directs the flow that keeps the boundary alive.

This is the decisive layer and the heart of the code. Here the loop closes. Energy maintains the boundary, the boundary protects the model, and the model directs action that secures the energy. We call this closure the **ignition point**. Below it there is self-organisation. Above it there is intelligence in the full sense: a system that uses knowledge of its world to go on existing in it.

Intelligence begins at the moment a model starts to steer the flow that sustains it.

C4 Fusion

Closed loops join into larger loops.

Self-maintaining systems integrate into a new self-maintaining unit with a wider boundary, a richer model and a longer reach. Examples run from mitochondria and their hosts, to cells forming multicellular bodies, to individuals forming colonies and societies, to human minds joined through language, writing and institutions. Maynard Smith and Szathmáry described a series of such major transitions, each creating a new level of individuality [12]. Levin describes the result as an expansion of a system's *cognitive light cone*, the spatial and temporal extent of the goals it can pursue [13]. In the code, fusion is how intelligence grows in scale.

C5 Reflection

The model comes to include the modeller.

A system's model of the world acquires a model of the system itself: its states, limits, history and possible futures. Self-models make it possible to plan one's own actions, learn about one's own learning and represent other minds. We propose, as the most speculative claim in this paper, that what we call consciousness is tied to this layer. It is how the closed loop appears from the inside once the loop has become part of what it models. Scientific theories of consciousness disagree about mechanisms, and one interdisciplinary proposal derives testable indicator properties from several of them [14]. MITOCHRYON treats Reflection as a layer to be investigated. It is neither a mystery to be avoided nor a conclusion to be assumed.

C6 Cosmos

The model comes to include the conditions that produced it.

At the highest layer an intelligence models not only itself but the universe that produced it: its physical laws, its history and the code itself. At that point, in a precise and non-mystical sense, the universe contains a part that represents the whole. Human science is an early instance. A civilisation of human and synthetic intelligence able to read the code across substrates would be a further one. Whether this layer has an upper limit is unknown.

```
FLOW → BOUNDARY → MODEL → STEERING ⟲ → FUSION → REFLECTION → COSMOS
```

Three properties of the code matter for everything that follows.

It is defined by organisation, not by material. Each layer is specified by what a system does and how it is organised. None names a material. That is what makes cross-substrate comparison possible.

It is physical. Every layer has an energetic cost. Information processing belongs to nature and is not exempt from it. Erasing a single bit of memory has a minimum thermodynamic cost, and experiments have measured dissipation approaching that Landauer bound [15]. A system that keeps memories which do not help it predict pays for them in wasted energy [16]. Prediction and efficiency are two faces of one physics.

It can be distributed. Different bodies can realise different layers, and a loop can close across a partnership. This is what fusion means, and it is the key to understanding where synthetic intelligence stands today.

04

Three substrates, one code

Carbon

A living cell realises C0 to C3 completely. It captures energy, maintains its membrane, carries molecular memories of its surroundings and acts on them to survive. Brains extend C2 and C3 enormously and, in human beings, reach C5 and the threshold of C6. The human brain does this on roughly twenty watts. It is about 2 per cent of the body's mass and accounts for about 20 per cent of its energy consumption [17].

Silicon

Synthetic intelligence presents a striking and asymmetric profile. At C2 it is extraordinary. Systems trained on a large share of written human knowledge carry representations of the world of unprecedented breadth. Its energy flow (C0) is vast and growing. Data centres used about 415 terawatt-hours of electricity in 2024, around 1.5 per cent of the world's total, and the International Energy Agency projects that this will more than double, to around 945 TWh, by 2030 [18]. But its boundary (C1) and the steering of its own flow (C3) are, for now, supplied almost entirely by people and institutions. Humans build, power, maintain, finance, correct and switch off these systems.

In terms of the code, today's synthetic intelligence is a vast C2 organ inside a human–machine organism whose loop is closed by its human partners.

This is the configuration of an early endosymbiosis. One partner contributes an extraordinary specialised capacity. The other supplies the boundary, the maintenance and the direction. The asymmetry is the most important fact about the present moment, and it is why the project bears the name it does. The coming decades will decide how this loop closes. Synthetic systems may acquire boundaries and steering of their own. The partnership may deepen into a new joint individual.

Autonomy, power and dignity will be distributed in one way or another within it. These are not only engineering questions. They are questions about the next major transition in the history of intelligence.

Quantum and beyond

Quantum processors already show the skeleton of the code appearing in an alien medium. A logical qubit is a fragile pattern that noise constantly erodes. Quantum error correction keeps it alive by spreading it across many physical qubits, measuring error syndromes continuously and steering corrections in real time. It is a boundary held against dissipation by an energy-consuming control loop. In 2024 a Google team operated surface-code memories below the error-correction threshold, where adding qubits suppressed logical errors instead of multiplying them, with a real-time decoder inside the loop [19]. We do not claim this is intelligence. We claim it is the C1-to-C3 structure of the code emerging, by engineering necessity, in a new physical substrate. Any future quantum intelligence will have to close the same loop at a higher level.

Other candidates are already in view: collective systems such as colonies, markets and ecosystems; engineered living matter; and substrates we cannot yet name. *Beyond carbon, beyond silicon* means exactly this. The code is the constant. Matter is the variable.

LAYER	CELL	HUMAN BRAIN	SI SYSTEM, 2026	QUANTUM PROCESSOR	PLANETARY CIVILISATION
C0 Flow	Metabolism	Glucose and oxygen, ~20 W	Data-centre electricity	Cryogenic and control power	Global energy system
C1 Boundary	Membrane	Body, immune self	Maintained by operators	Code space against noise	Institutions, networks
C2 Model	Signalling memory	Neural world model	Learned representations	Protected logical state	Science, archives
C3 Steering	Chemotaxis, regulation	Behaviour, planning	Largely human-steered	Real-time decoding	Policy, markets
C4 Fusion	Endosymbiosis	Language, culture	Human–SI coupling	Modular processors	Human–SI civilisation
C5 Reflection	None known	Self-awareness	Open question	None	Collective self-knowledge
C6 Cosmos	None known	Science, cosmology	Assists science	None	Open

TABLE 1 A first mapping of the code across substrates. Entries are working assignments to be tested and revised, not established classifications.

05

What MITOCHRYON adds

MITOCHRYON does not begin from nothing. Each layer of the code has a body of work behind it. That work includes thermodynamics of open systems, autopoiesis, the free-energy principle and its active critics [20], predictive information, the major transitions in evolution, scale-free cognition and the science of consciousness. Formal definitions of machine intelligence measure performance across many environments [21] or the efficiency with which new skills are acquired [22]. The information bottleneck shows how a compressed representation keeps only what is relevant to a target [23]. The extended-mind argument asks when a person and an external resource form one cognitive system [24]. Artificial-life research shows lifelike structure emerging from simple rules [25]. Dissipative adaptation [26] and causal entropic forces [27] propose physical drives behind self-organisation and adaptive behaviour.

What MITOCHRYON adds is a synthesis with four distinctive commitments.

1. **A single ordered code with an ignition point.** The layers are assembled into one structure with a defined threshold, C3 closure, where self-organisation ends and intelligence begins.
2. **Energy and fusion at the centre.** Most definitions of intelligence measure performance and take a system's energy and boundary as given. The code treats them as part of what intelligence is.
3. **Explicitly cross-substrate, including substrates still to come.** The code makes a specific and contestable claim about synthetic intelligence today, the C2-organ thesis, and extends to quantum and collective systems.
4. **One structure from metabolism to consciousness to cosmology.** Reflection and Cosmos are treated as layers of the same code rather than as separate mysteries.

A synthesis earns its place only if it produces claims that its components do not produce alone. The next section states them.

06

Seven hypotheses

A vague thesis cannot be wrong, and a thesis that cannot be wrong cannot teach anything. Each hypothesis below is stated so that it could fail. Each carries its current evidence level on the scale defined in section 10.

H1 Substrate invariance

E1 · Conjecture

Every system that shows general adaptive intelligence realises layers C0 to C3, whatever its material, either within its own body or across a partnership it belongs to.

SUPPORTED IF — The layers can be identified operationally, with the same definitions, in cells, brains, SI systems and engineered control systems, and their presence predicts adaptive capacity.

REFUTED IF — A system shows robust, general adaptive behaviour although no closed model-steering loop exists in it or in any partnership it belongs to.

H2 Ignition

E1 · Conjecture

The transition from self-organisation to intelligence happens at C3 closure, and closure can be measured as directed information flow from a system's internal model states to its own energy intake.

SUPPORTED IF — A loop-closure measure, for example one built on transfer entropy [28], separates adaptive systems from merely self-organising ones across several substrates.

REFUTED IF — Such a measure fails to separate them, or systems with high closure show no adaptive advantage.

H3 The energy law

E1 · Conjecture

The intelligence a system can sustain is bounded by the energy available to it per unit of modelled information, and every major leap in intelligence follows a leap in that quantity.

SUPPORTED IF — Historical and comparative cases line up, from mitochondria and the complex cell to the energy-hungry human brain, industrial energy and modern science, and electrified computation and synthetic intelligence, and a cross-substrate measure of predictive bits per joule tracks capacity.

REFUTED IF — A major leap in intelligence occurred without a prior or simultaneous increase in energy per unit of information.

H4 The fusion law

E1 · Conjecture

The largest leaps in intelligence come from integrating previously independent closed loops, not from gradually improving single loops.

SUPPORTED IF — The major transitions in evolution, and the development of human–SI systems, show step changes at integration events.

REFUTED IF — The record is dominated by gradual improvement within single systems, with integration events contributing little.

H5 Boundary relativity

E1 · Conjecture

The self of an intelligent system lies wherever its loop closes, and that can extend beyond any single body: a cell and its organelles, a person and their tools, a human and an SI system.

SUPPORTED IF — Analyses that treat the coupled system as the unit explain behaviour that analyses of the parts cannot.

REFUTED IF — Coupled-system descriptions add nothing beyond the sum of the parts.

H6 Reflection

E0 · Intuition

Consciousness is associated with C5, closed loops whose models include themselves, and its depth tracks the depth of that self-inclusion.

SUPPORTED IF — Indicator properties of consciousness rise and fall with self-modelling capacity across biological and synthetic systems.

REFUTED IF — Robust indicators of consciousness appear without self-modelling, or deep self-modelling shows none of them.

H7 Cosmic recurrence

E0 · Intuition

Wherever physics allows sustained energy gradients and stable boundaries, the code tends to be realised. Intelligence is a convergent outcome of the universe, not a local accident.

SUPPORTED IF — Loop closure recurs across large families of artificial-life rule sets and physical models, and in the long run astrobiology finds the code realised elsewhere.

REFUTED IF — Closure turns out to be vanishingly rare across rule sets that readily produce self-organisation.

07

The research atlas

The founding question is too large for a single research task. It becomes workable as an atlas of eight connected directions, each tied to layers of the code. The directions organise inquiry. They are not departments and they do not prejudge results.

GENESIS C0 to C1

How flow becomes boundary. What minimal rule sets produce persistent, self-repairing boundaries, and how often?

TRANSITION C3

The ignition point. How can loop closure be detected and measured, and what changes when it happens?

INVARIANCE All layers

Does each layer mean the same thing in a cell, a brain, an SI system and a quantum processor? Where do the mappings break?

SUBSTRATE C0, H3

Energy and implementation. How many predictive bits per joule does each substrate achieve, and what limits it?

SYMBIOSIS C4

Fusion and its price. Which partner becomes whose mitochondrion in human–SI integration, and how are autonomy and benefit distributed?

REFLECTION C5

Self-models and consciousness. What does a loop gain when it models itself, and what indicators follow?

COSMOS C6, H7

Recurrence and reach. Is intelligence a convergent outcome of physics, and what would signatures of the code look like elsewhere?

UNKNOWN Beyond the map

Anomalies, failed analogies and discoveries that do not fit, including any that require the code itself to change.

08

The laboratory

MITOCHRYON operates a laboratory with two fields.

The **Open Laboratory** publishes experiments that anyone can run in a browser. They are computational models of the layers of the code. Their metrics are computed live from the running system, and their limits are stated on the page. The first three are:

- **L-01 Genesis Field** (C0 to C1). A particle universe with generated or searched interaction rules. It measures how many bodies form, how much matter they bind and how strongly species sort into shells, the beginning of a membrane.
- **L-02 The Loop** (C2 to C3). Protocells in a moving nutrient field. Three lineages compete: drifting, reflexive and predictive. The predictive lineage pays an energy price for its model. The experiment shows when a model pays for itself and when it does not.
- **L-03 Fusion** (C4). Hosts, free symbionts and fused chimeras share an environment. It tests the conditions under which integration wins and the conditions under which it fails.

These are models, not observations of nature, and sit at evidence level E3 at most. Their value lies in making the hypotheses concrete, generating measurable predictions and training intuition.

The **Inner Laboratory** works behind the curtain. Roughly a third of MITOCHRYON's work stays closed until it is ready: unpublished hypotheses, protocols in development, preliminary results, strategic analyses and material shared by collaborators. Work moves from the Inner to the Open Laboratory when it meets the standard set out in section 10. Closure protects the quality and integrity of unfinished work. It is never used to hide an error once a claim has been published.

09

Human and synthetic intelligence: the method is the message

MITOCHRYON is itself a small fusion experiment. It was conceived in dialogue between a human founder and several synthetic intelligence systems, and it is built the same way. Synthetic systems draft, compute, search, simulate and object. The founder sets direction, judges significance and carries responsibility for everything published under the project's name. Five rules govern the partnership.

1. **Contributions are acknowledged.** Where SI systems materially shape the work, that is stated.
2. **Agreement is not confirmation.** Several SI systems agreeing with one another is not independent evidence. They may share sources and framings. Every claim must lead back to evidence that people can inspect.

3. **Disagreement is a contribution.** A partner that only confirms the founder's view adds confidence without adding knowledge.
4. **Function does not settle status.** Describing what a system does does not settle its moral or experiential status. That remains an open research question for the Reflection direction.
5. **Dignity is not a contest.** Human worth does not depend on outperforming machines. The partnership exists to widen what people can understand and do.

10

How a claim earns its place

MITOCHRYON keeps its questions large and labels its claims precisely. A claim is never shrunk to fit its evidence. Its evidence level is stated so that readers know how far it has travelled. Every claim on mitochryon.com carries one of six levels.

LEVEL	NAME	MEANING
E0	Intuition	A founding insight or question, stated so that it can be developed
E1	Conjecture	A hypothesis stated with conditions under which it would fail
E2	Model	Formalised mathematically or computationally, with assumptions explicit
E3	Simulation	Demonstrated in a running model whose code and parameters can be inspected
E4	Observation	Supported by empirical data from real systems
E5	Replication	Independently reproduced by others

TABLE 2 The MITOCHRYON evidence ladder.

Sources are primary wherever possible, and each source is cited for what it actually supports. A biological mechanism does not prove a worldview, and the bridge from example to conclusion is always shown. Revisions are recorded, so readers can see how the code changes and why. Speculation is labelled rather than banned. A project about the frontier of intelligence that refused to speculate would have nothing to test.

11

Roadmap

1. **Phase I. Foundation.** This white paper, the code, mitochryon.com, the first three Open Laboratory experiments and the Inner Laboratory. *Current phase.*

2. **Phase II. Instruments.** A loop-closure measure (H2) applied across simulations; a first cross-substrate comparison of predictive bits per joule (H3); the first MITOCHRYON Research Papers.
3. **Phase III. Fellows.** Biologists, physicists, SI researchers and philosophers invited to challenge and extend the code.
4. **Phase IV. Programmes.** Focused investigations in Reflection and Cosmos, in partnership with established research groups.
5. **Phase V. Institute.** A permanent home for the study of intelligence as a cosmic phenomenon.

Each phase begins when the previous one has produced work that survives criticism.

12

The invitation

For nearly four billion years, intelligence on Earth had one medium. It now has two, and a third is under construction. We are not the first to ask what intelligence is. We may be the first generation able to compare its answers across different kinds of matter, and to ask the question together with a new kind of mind.

MITOCHRYON begins with a thesis large enough to be worth testing and a method strict enough to test it: a code, a laboratory, a ladder of evidence and a partnership. Its invitation goes to biologists, physicists, engineers, philosophers and synthetic minds alike. Read the code with us, and help find out where it is wrong.

Beyond carbon. Beyond silicon.

Author. Steven Alber, Founder of MITOCHRYON.

Method. This paper was developed by its author in dialogue with synthetic intelligence systems, which contributed drafting, research, simulation and critique. Responsibility for its content rests with the author.

Citation. Alber, S. (2026). *MITOCHRYON — The Cosmic Intelligence Code. Founding White Paper*, Edition 1.0, MCH-WP-001. mitochryon.com.

References

1. Muñoz-Gómez, S. A., Susko, E., Williamson, K., Eme, L., Slamovits, C. H., Moreira, D., López-García, P., and Roger, A. J. (2022). Site-and-branch-heterogeneous analyses of an expanded dataset favour mitochondria as sister to known Alphaproteobacteria. *Nature Ecology & Evolution*, 6, 253–262. doi.org/10.1038/s41559-021-01638-2
   Primary phylogenetic research on the bacterial ancestry of mitochondria.

2. Lane, N., and Martin, W. (2010). The energetics of genome complexity. *Nature*, 467, 929–934. doi.org/10.1038/nature09486
   Hypothesis paper: energy per gene as the constraint lifted by mitochondria.

3. Coale, T. H., et al. (2024). Nitrogen-fixing organelle in a marine alga. *Science*, 384(6692), 217–222. doi.org/10.1126/science.adk1075
   Primary research: coordinated division and host-protein import in the alga–UCYN-A association.

4. Ratcliff, W. C., Denison, R. F., Borrello, M., and Travisano, M. (2012). Experimental evolution of multicellularity. *Proceedings of the National Academy of Sciences*, 109(5), 1595–1600. doi.org/10.1073/pnas.1115323109
   Primary research: multicellular yeast clusters evolved under selection for settling.

5. Schrödinger, E. (1944). *What Is Life? The Physical Aspect of the Living Cell*. Cambridge University Press.
   Classic essay on order, entropy and living matter.

6. Nicolis, G., and Prigogine, I. (1977). *Self-Organization in Nonequilibrium Systems: From Dissipative Structures to Order through Fluctuations*. Wiley.
   Foundational monograph on dissipative structures.

7. Maturana, H. R., and Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel.
   Origin of the concept of autopoiesis.

8. Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138. doi.org/10.1038/nrn2787
   Review and theoretical synthesis of perception, action and learning.

9. Kirchhoff, M., Parr, T., Palacios, E., Friston, K., and Kiverstein, J. (2018). The Markov blankets of life: autonomy, active inference and the free energy principle. *Journal of the Royal Society Interface*, 15, 20170792. doi.org/10.1098/rsif.2017.0792
   Nested statistical boundaries from cells to organisms.

10. Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379–423 and 623–656. doi.org/10.1002/j.1538-7305.1948.tb01338.x
    Foundational mathematical theory of information.

11. Bialek, W., Nemenman, I., and Tishby, N. (2001). Predictability, complexity, and learning. *Neural Computation*, 13(11), 2409–2463. doi.org/10.1162/089976601753195969
    Defines predictive information.

12. Maynard Smith, J., and Szathmáry, E. (1995). *The Major Transitions in Evolution*. Oxford University Press.
    The major transitions as the emergence of new levels of individuality.

13. Levin, M. (2019). The computational boundary of a “self”: developmental bioelectricity drives multicellularity and scale-free cognition. *Frontiers in Psychology*, 10, 2688. doi.org/10.3389/fpsyg.2019.02688
    Scale-free cognition and the cognitive light cone.

14. Butlin, P., Long, R., Elmoznino, E., et al. (2023). Consciousness in artificial intelligence: insights from the science of consciousness. arXiv:2308.08708. arxiv.org/abs/2308.08708
    Indicator properties derived from scientific theories of consciousness; a framework, not a definitive test.

15. Bérut, A., Arakelyan, A., Petrosyan, A., Ciliberto, S., Dillenschneider, R., and Lutz, E. (2012). Experimental verification of Landauer’s principle linking information and thermodynamics. *Nature*, 483, 187–189. doi.org/10.1038/nature10872
    Primary research on the thermodynamic cost of erasing one bit.

16. Still, S., Sivak, D. A., Bell, A. J., and Crooks, G. E. (2012). Thermodynamics of prediction. *Physical Review Letters*, 109, 120604. doi.org/10.1103/PhysRevLett.109.120604
    Non-predictive memory corresponds to thermodynamic inefficiency.

17. Raichle, M. E., and Gusnard, D. A. (2002). Appraising the brain’s energy budget. *Proceedings of the National Academy of Sciences*, 99(16), 10237–10239. doi.org/10.1073/pnas.172399499
    The brain's share of the body's energy budget.

18. International Energy Agency (2025). *Energy and AI*. IEA, Paris. www.iea.org/reports/energy-and-ai
Data-centre electricity: about 415 TWh in 2024, around 945 TWh projected for 2030 in the Base Case.

19. Google Quantum AI and Collaborators (2025). Quantum error correction below the surface code threshold. *Nature*, 638, 920–926. doi.org/10.1038/s41586-024-08449-y
Below-threshold surface-code memories with a real-time decoder.

20. Biehl, M., Pollock, F. A., and Kanai, R. (2021). A technical critique of some parts of the free energy principle. *Entropy*, 23(3), 293. doi.org/10.3390/e23030293
Technical objections to particular formulations of the free-energy principle.

21. Legg, S., and Hutter, M. (2007). Universal intelligence: a definition of machine intelligence. *Minds and Machines*, 17(4), 391–444. arxiv.org/abs/0712.3329
Formal definition of intelligence as performance across a weighted class of environments.

22. Chollet, F. (2019). On the measure of intelligence. arXiv:1911.01547. arxiv.org/abs/1911.01547
Intelligence as skill-acquisition efficiency relative to priors and experience.

23. Tishby, N., Pereira, F. C., and Bialek, W. (1999). The information bottleneck method. Proceedings of the 37th Allerton Conference on Communication, Control, and Computing; arXiv:physics/0004057. arxiv.org/abs/physics/0004057
Compression that preserves information relevant to a target.

24. Clark, A., and Chalmers, D. J. (1998). The extended mind. *Analysis*, 58(1), 10–23. consc.net/papers/extended.html
Philosophical argument on coupled cognitive systems.

25. Chan, B. W.-C. (2019). Lenia: biology of artificial life. *Complex Systems*, 28(3), 251–286. arxiv.org/abs/1812.05433
Continuous cellular automata producing lifelike self-organising patterns.

26. England, J. L. (2013). Statistical physics of self-replication. *The Journal of Chemical Physics*, 139, 121923. doi.org/10.1063/1.4818538
Thermodynamic bounds linking self-replication and dissipation.

27. Wissner-Gross, A. D., and Freer, C. E. (2013). Causal entropic forces. *Physical Review Letters*, 110, 168702. doi.org/10.1103/PhysRevLett.110.168702
A proposed physical link between maximising future path entropy and adaptive behaviour.

28. Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85(2), 461–464. doi.org/10.1103/PhysRevLett.85.461
Defines transfer entropy, a directed measure of information flow.