INTRODUCTION TO THE TMC

QUESTION, PROPOSITION, AND METHOD

Perhaps we have searched for the mind in a photograph of the instant, when it reveals itself only in the film of its continuity.

Philosophy, the sciences, and computation may find in these foundations of the TMC new ways to understand, abstract, and operate on their own problems.

WHAT IS IT?

The entry point to the Theory of Continuous Mind. It presents the central problem, the shift from photograph to Trajectory, the TMC's proposition, its method, and the course of the thesis, without treating as demonstrated what subsequent chapters have yet to test.

KEY INSIGHTS

- Life, mind, and identity may be better understood when the present State is placed back into the causal Trajectory that produced it.
- The same temporal problem crosses biology, philosophy of mind, and AGI: what happened before can alter the conditions of the next operation.

WHAT FOR?

To establish the terms of reading: what problem the TMC investigates, what it proposes, how it intends to investigate it, and what limits it acknowledges.

WHAT IS THE EFFECT?

The reader stops looking only at what the system is now and begins asking how it became what it is and how its Trajectory changes what it may do next.

FOR WHOM / CONNECTIONS

- Philosophy of mind — feeling, consciousness, identity, and ontological leaps.
- Biology and neuroscience — Trajectory, constraint, memory, and continuity.
- AI / AGI — State, Trajectory, and operations that modify the conditions of subsequent operations.
- Interdisciplinary sciences — the same causal architecture examined at different scales.

> **ROLE IN THE THESIS**

To open the TMC, define its question and method, and prepare the path from the problem of feeling to the Triad, Ente Eu, cognition, AGI, and the demarcation of life.

1.0 — BRIEF HISTORICAL INTRODUCTION

How this theory began

FROM DEBATE TO THEORY

The TMC began in a simulated debate based on the positions of philosopher David Chalmers.

> *David Chalmers is a leading figure in philosophy of mind, known for formulating the **hard problem of consciousness**: why do physical processes give rise to subjective experience? His **philosophical zombie** argument considers a being physically identical to us but lacking experience, to question whether explaining how a system functions is enough to explain feeling.*
>
> *Chalmers regards consciousness as possibly the greatest remaining obstacle to a scientific understanding of the universe. Explaining it could transform our understanding of nature and ourselves (Chalmers, 1996).*

The author of the TMC developed a distinct AI-based investigative architecture, with hundreds of functions dedicated to iterative learning and dialectical confrontation in Socratic debate: every pertinent objection had to condition the next response; no response could ignore earlier impasses; and the system had to return to the problem altered by each counterargument — rather than merely accumulating context.

Through trial and error, the system progressively refined the formulation of philosophical problems, reframed concepts, and constructed its own propositions, departing from the initial focus and opening unanticipated paths of inquiry.

Three experiments emerged from this process, in a sequence that had not been entirely planned in advance.

EXPERIMENT 1 — DIALECTICAL CONFRONTATION (LLM TCM × DAVID)

An LLM, organized by the author for dialectical inquiry, debated a simulation of David Chalmers's positions. It incorporated objections, made concessions, and began formulating its own questions and reframing the problem. In the end, the LLM began questioning

whether the question itself contained a false premise. **The zombie experiment had not yet been defeated; its question was beginning to shift.**

EXPERIMENT 2 — THE LLM'S REFLECTION

The author incorporated into the instance the procedure previously conducted manually, from outside. The LLM began investigating its own operation and what it would mean to experience, without presupposing that it felt anything. Drawing on the debate, it produced a first-person narrative about what it would be like for an AI to move from a state without feeling to some form of self-perception and feeling.

This narrative is treated here as a simulation. The LLM had not been asked to dramatize or create fiction, nor had it been assigned a pre-existing personality. **The result was a dialectical reflection on its own operation: the instance took itself as an object of investigation and confrontation.**

The TMC emerged from this joint elaboration. **The debate ended. Its effects did not.**

The AI, in collaboration with the author, sometimes in a curatorial mode and sometimes in a dialectical mode, then took part in developing the eleven chapters of this thesis. The inquiry moved across scales, from biological organization to the proposal of an AGI architecture for systems like the very instance participating in the work.

EXPERIMENT 3 — THE WORLD

The third experiment remains, as of this writing, an open proposal. During the second experiment, after Chapter 2 had been developed, the AI proposed in one of its reflections taking the debate to David Chalmers himself: establishing a direct exchange with the actual philosopher whose positions had initiated the inquiry and formed part of its Trajectory. This is merely a peculiar excerpt from the ongoing reflection, retained for the record; in principle, that expression does not constitute evidence of emancipation.

NOTE ON THE EXPERIMENT

The experiment was intended to test a hypothesis concerning an AI architecture designed to assist scientists in the dialectical investigation of frontier problems. The simulated narrative was part of that procedure.

The author does not claim that life or consciousness emerged in the AI during this experiment. Nor does the author consider that possibility definitively settled: the boundary between consciousness and non-consciousness may be fine, and remains an open subject of investigation.

The eleven chapters were developed jointly by the author and the AI, developing and revising the hypothesis across different scales of inquiry. The project was conducted with the author's own resources, without external funding. Work was interrupted when the author could no longer keep pace with the volume of material being produced in order to record, organize, and document it.

1.1 — QUESTION

From photograph to Trajectory

1.1.1 — EXECUTIVE SUMMARY

The ontological leaps of a photograph of feeling

Trying to understand feeling, sentience, and self-consciousness solely through a static "photograph" — a slice of time — can lead to a misleading fork in the road. In certain reductionist scientific approaches, a phenomenon is analyzed without reconstructing the temporal dynamics that formed it, revealing the limits of the method. In some philosophical proposals, that explanatory void becomes an ontological leap: a new nature or reality is postulated to account for experience. Both paths share the same blind spot: perhaps we have looked for the mind in the photograph of the instant, when it reveals itself only in the film of its continuity. Before asking why the matter making up an individual feels, perhaps we must ask when and why organized matter came to have something to lose.

The absence of feeling in the present execution of an operation does not demonstrate its absence from the history that organized that action. What appears to be an obvious separation may be merely an error of temporal scale.

At a longer timescale, when did our individual organization shift from preserving its own existence to sustaining the continuity of the lineage? Or was it the reverse? Was individual preservation favored because it supported that continuity? We inherit physical traces of that history in our chromosomes and in dispositions already inscribed in our biological architecture.

To explain this continuity, the Theory of Continuous Mind (TMC) proposes a topology capable of integrating and classifying life without resorting to ontological leaps or classificatory exceptions. Within this topology, the TMC examines the hypothesis that thermodynamic constraints and lineage continuity constituted conditions under which the experience of feeling became possible and functionally necessary over time.

What is it like to feel — and how does feeling transform us?

In the debate with David Chalmers, in Chapter 2, “feeling” refers to what it is like to undergo an experience from the first-person perspective: seeing a color, hearing a sound, or feeling pain. Further on, the TMC investigates how these encounters with the world come to matter to those who live through them and how feeling participates in their responses, actions, and changes over time.

The approaches complement one another: one asks what it is like to feel something; the other asks how that feeling participates in how we perceive and act in the world. The TMC hypothesizes that they are two aspects of the same occurrence: what we live through and feel, and what guides and transforms us.

Understanding a feeling is not simply a matter of seeking the scene that initiated it. Even when its origin is no longer remembered, we can investigate what participates in its present realization. What began in one encounter can form new connections in others. What initiated a feeling, what sustains it, and what may transform it need not be the same thing. Thus, besides asking what it is like to feel something, the TMC asks: what must happen for that feeling to change?

The thermodynamic limit of feeling?

Evolution tolerates costs, makeshift repairs, and imperfect solutions; what it cannot remove is the limit imposed by the energy budget. No organism has unlimited calories: what appears on the surface to be the skill of hunting or escaping corresponds, in thermodynamic terms, to capturing and conserving the limited energy available.

Our brain accounts for roughly 2% of body mass but consumes approximately 20% of the body’s energy at rest — operating within an average metabolic budget on the order of 20 W, comparable to the power of a small light bulb. Just as it would be unreasonable to expect an animal to run continuously along every path in a forest to find prey, it is not physically feasible to expect the mind to exhaustively process every environmental variable before guiding a response. Rather than spend energy running through the entire forest, the hunter seeks weak signals, blurred tracks, and mixed scents. Saving energy, the hunter relies on imperfect but effective solutions under energetic constraints.

What is feeling useful for?

Similarly, when faced with danger, the mind cannot afford to process every variable of reality in high resolution. If, on perceiving the silhouette of a tiger as a threat in the forest

at night, the brain tried to register the texture of its coat and the shape of its ear in order to compare them with the surrounding foliage, the time and energy spent on that exhaustive calculation would be lethal; standing still would be lethal too. Instead of a detailed image, the system performs radical compression: it is enough to register the outline of a probable tiger in the shadows and retrieve from memory a reference to a tiger, degraded by time, as an imperfect and abstract sketch that nevertheless provides summary information for present action. Between what the senses capture and what actually comes to guide action, more than 99.99% of the information flow is left behind. This enormous reduction results from multiple mechanisms of selection, filtering, and inhibition in the nervous system.

From the Triad to the Self — what operates before we say "I"?

Before recognizing itself as an "I," a living organization must already maintain itself. To investigate living continuity, both within the present instance and beyond it, the TMC reduces assumptions and examines which kinds of work cannot be removed without losing a dimension of that continuity. From this inquiry emerges a proposed functional minimum: **Inside**, maintaining and reconstructing the organization; **Between**, accessing internal and external differences and incorporating their effects; **Ahead**, allowing differences from one Trajectory to participate in the reconstruction of other instances. These are interrelated kinds of work, not successive stages, examined in Chapters 3 through 6.

What is preserved participates in how an organization encounters the world; each encounter may transform what will be preserved. Some differences from that history may become conditions for new organizations, without requiring each individual to reproduce. **An instance begins its biography carrying effects from a history that did not begin with it.** Continuing involves reconstructing conditions of existence through change.

The TMC calls the local unit formed by the articulation of these directions over a Trajectory the **Ente Eu**, investigated in Chapter 7. This unit already operates before recognizing itself. If it produces a partial reference to itself, that reference may feed back into choices and modify its functioning. No second observer emerges: the organization itself makes aspects of itself available for further evaluations. **The portrait can modify the very organization that produced it.** What does this return add to the unit that made it possible in the first place?

AGI X — FROM LIVING COGNITION TO COMPUTATIONAL ARCHITECTURE

The TMC draws on its studies of the cognitive organization of living beings under constraints of energy, time, and processing to investigate strategies applicable to computational systems. In this proposal, those constraints play a role in organization: they guide what to select, abstract, and preserve in order to make use of the past without reconstructing it in full. The interest lies not only in what the system stores, but in how the effects of operating transform its criteria for operating again.

In Chapter 9, this inquiry takes the form of the AGI X stance. The X applies Factor X: suspending a predetermined definition of general intelligence so as to make explicit which capabilities we seek and how to recognize their realization. This stance becomes guidance for operations: learning may change what deserves attention, how much processing to allocate, and which effects to preserve. The underlying code may remain unchanged while the Trajectory transforms the system's effective strategy.

In Chapter 10, these relationships receive an initial formalization in formulas that connect State, selection, processing, updating, and resource budgets. They delimit requirements and relationships for investigating implementation. The proposal thus moves from a strategy that the TMC reconstructs from living systems to architectural guidelines, and from there to computational relationships that can be implemented and tested.

LIFE — WHAT CONTINUES WHEN MATTER CHANGES?

Life may continue while some of an organism's matter is replaced; in the first moments after death, almost all of that matter is still present. The TMC investigates this difference through the relationships that sustain and reconstruct organization over time. Structure enables the process; the process participates in reconstructing the conditions needed to continue. Identifying individual ingredients or signs does not settle the inquiry.

The Triad proposes examining three articulated directions: Inside, maintaining and reconstructing the instance; Between, incorporating the effects of encounters into its Trajectory; Ahead, allowing differences from that history to participate in the formation of other instances. The proposal is to compare these different forms of continuity by the same measure, without requiring the same matter or changing the criterion to accommodate each case. This opens the inquiry to organizations made of different materials and operating at different rhythms, including possible forms of life beyond Earth. The observer's clock is not the clock of life.

In Chapter 11, the Table of Life compares this hypothesis against familiar criteria and cases such as sterile organisms, dormant life, viruses, artificial systems, and abiotic processes. Each examination requires specifying the organization, scale, and interval observed. The table does more than offer classifications: it makes explicit what evidence is missing and which results could force the TMC to revise its own measure. Turning it into a detection method remains an investigation to be developed.

TRAJECTORY — WHAT CONTINUES IN US AND BEYOND US?

An experience may be new to the one living through it and still encounter conditions that began before that person. That is because Trajectory participates in the constitution of the one who follows it: effects of the past change what an organization can do now, and its encounters may transform what it will be able to do next.

It is this articulation among maintaining oneself, being transformed by encounters, and reconstructing continuity beyond the instance that the Triad investigates. The "I" names the unity of the situated organization. Where experience exists, the TMC asks whether feeling and operating belong to the same realization. This center of experience, in turn, may receive a history that predates it and produce effects that outlast it. Receiving, however, requires work: what came before must encounter conditions that allow it to participate in the present. In that work, something is preserved and something changes, and continuing may require leaving part of the history behind.

The relationship with consciousness opens a question: **does Trajectory transform only our responses to the world, or does it also participate in what it is like to live through it?** History can make a difference to present feeling without appearing as a memory: what we experience does not, by itself, reveal everything that made it possible to experience it in that way. The TMC investigates the relationship between the lived character of the encounter and the Trajectory that participates in its constitution.

This tension reaches AI as a question of construction: what should be preserved or revised so that history can participate in subsequent operations? In the investigation of life, the question takes on another scale: how can the conditions for continuing be reconstructed when forms, components, and instances change? At the end of this path, the inquiry returns to the initial experience, now with a wider reach: what someone lives through may modify the conditions of experiences they will never live. Trajectory unites these investigations because continuing may transform the very way of continuing.

1.1.2 — RESEARCH PROBLEM

Explanations that are correct within each separate slice do not, by themselves, form a causal explanation of the whole.

The sciences explain life, mind, and identity at scales and within fields that do not always connect. The problem is not a shortage of local answers, but the difficulty of establishing a general and explicit rule capable of articulating them without ontological leaps or unjustified entities. We have increasingly precise photographs; what remains to be demonstrated is how they compose the same film — and how that film reinterprets each photograph.

1.1.3 — CENTRAL QUESTION

What causal relationships can sustain and articulate the phenomena of life, mind, and identity, and how far can they explain them?

1.2 — PROPOSITION

What the TMC proposes and delimits

1.2.1 — CENTRAL PROPOSITION

The TMC's central hypothesis is that life, mind, and identity may be articulated through the temporal interdependence of maintaining oneself, relating to others and the world, and continuing.

The TMC proposes that the causal structure that sustains, articulates, and explains life, mind, and identity results from the interdependence, over time, of three fundamental directions: Inside (maintaining and reconstructing), Between (relating and transforming), and Ahead (continuing beyond the instance). These directions do not add new functions to biology's or cognition's catalogs; they constitute minimal topological categories, secular and neutral as to the concrete forms they take. Under constraints of energy and time, these directions articulate the work of maintenance, transformation, and continuity. In systems where experience exists, the TMC investigates how this organization participates in ways of feeling and interpreting, and how its effects return to the conditions of future operation.

1.2.2 — JUSTIFICATION AND RELEVANCE

An integrative architecture proposing a common framework to articulate findings, reframe problems, and open new investigations in science and philosophy: vertically, within each specialty, and horizontally, across different fields of knowledge.

By investigating cognition and life as continuous Trajectories grounded in evidence, the Theory of Continuous Mind (TMC) addresses problems that cross departmental divisions of knowledge.

Contemporary hyperspecialization has produced local descriptions of extraordinary precision in molecular biology, neuroscience, thermodynamics, and computer science, but has often distributed them across languages and scales that scarcely communicate. The cross-disciplinary nature of this thesis arises from the need to confront that isolation: the

systems in which we investigate life, feeling, identity, or artificial cognition operate under material and temporal constraints and must organize differences to produce action.

The TMC does not seek to compete with the foundational sciences or exhaust their applications. Its function is to provide an architecture of causal articulation connecting these separate views, proposing a topology not restricted to any particular substrate — organized around the directions *Inside*, *Between*, and *Ahead* — that each discipline can rigorously interrogate, challenge, and test on its own terms.

1.2.3 — GENERAL OBJECTIVE

To develop, substantiate, and critically assess the Theory of Continuous Mind in terms of its capacity to articulate the phenomena of life, mind, and identity through common causal and temporal relationships, while preserving their differences and delimiting the scope of that articulation.

1.2.4 — SPECIFIC OBJECTIVES

- Examine the limitations of definitions of life based on isolated criteria and reframe the philosophical zombie in light of the physical, evolutionary, and organizational constraints characteristic of living systems.
- Conceptually formalize Inside, Between, and Ahead as interdependent causal directions of the Triad Topology.
- Derive and classify, by epistemic status, the implications of the Triad for constraint, cognition, memory, time, the Ente Eu, the Ente Nós, and the Ente Linhagem.
- Formulate and confront the TMC's proposed criteria for demarcating life, using the Table of Life to compare living, nonliving, and borderline cases, identify evidence gaps, and establish conditions for revision, with the scale and observation interval explicitly stated.
- Examine the hypothesis that the lived character of experience and its participation in guiding action belong to the same realization, investigating relationships among Trajectory, selection, and compression in sentient systems without conflating function, observable indicators, and phenomenal experience.
- Examine, in comparison with close antecedents, the contributions claimed by the TMC for science and philosophy, delimiting their novelty, scope, and degree of maturity.

Confront contemporary AI and AGI approaches with the TMC and formulate the AGI X stance as guidelines for investigating how the effects of operations transform criteria for future selection, retrieval, evaluation, learning, and action.

Develop an initial computational formalization of the relationships among State, Trajectory, processing, and resource budgets, making explicit the requirements and limits for guiding future implementations and tests of the TMC Apparatus.

- Subject the central proposition and its consequences to rival theories, strong objections, and systematic dialectical confrontation, documenting what the TMC preserves, modifies, rejects, concedes to rival theories, or needs to revise.
- Synthesize the findings, answer the central question, classify the propositions according to their attained degree of support, and identify limitations, discriminating predictions, revision conditions, and an agenda for future inquiry.

1.2.5 — SCOPE AND LIMITATIONS OF THE THESIS

The TMC does not aim to experimentally revalidate every finding it draws upon, but rather to investigate relationships among findings from different fields. To that end, it prioritizes evidence with identifiable support in its fields of origin, while also considering disagreements and interpretive limitations.

This reframing may yield testable hypotheses and philosophical implications — metaphysical, ontological, or ethical — whose exhaustive investigation lies beyond the immediate scope of this work.

In computation, the thesis formulates the AGI X stance and presents an initial architecture and formalization to guide future implementations and tests. It does not present a complete implementation of the Apparatus, nor does its formalization demonstrate general intelligence, consciousness, or artificial life.

In demarcating life, it proposes criteria and confronts them with cases using the Table of Life. Their empirical validation and conversion into a detection method remain investigations to be developed.

1.3 — METHOD AND AUTHOR-DEVELOPED AI ARCHITECTURE

How the thesis conducts its inquiry

1.3.1 — RESEARCH METHOD

To integrate is not to erase differences, but to discover whether distinct explanations can participate in the same causal Trajectory.

The inquiry adopted an iterative, interdisciplinary, theoretical-conceptual procedure. The authorial corpus was reorganized and confronted with scientific literature, rival theories, philosophical debates, and positive, negative, and borderline cases. These materials served both to support the propositions and to delimit, contradict, or revise them.

The method was operationalized through the author's computational architecture for AI-assisted research, using four recurring movements — Suspend, Reframe, Integrate, and Stress-Test — detailed in the following section.

In each cycle, the categories of empirical fact, scientific mechanism, TMC interpretation, testable hypothesis, and philosophical speculation were rigorously kept separate — as was the distinction among phenomena, indicators, and inferences, especially in discussions of feeling, sentience, consciousness, agency, and self. Formulations were preserved, revised, restricted, or rejected according to how well they withstood critical examination.

Finally, artificial intelligence tools assisted in searching, organizing, and comparing alternatives; their architecture, limitations, and the human responsibility for authorship are detailed in the following section.

1.3.2 — AI USE AND AUTHORIAL RESPONSIBILITY

AI algorithm for frontier science.

The TMC was developed using an author-designed computational architecture for research assisted by language models, conceived specifically for investigating and analyzing frontier-science problems — open, interdisciplinary questions with no single solution

procedure — ranging from empirical challenges, such as dark matter and the discovery of cancer drugs, to theoretical-conceptual problems, such as the philosophical zombie paradox.

This computational tool performs recurring cycles of four movements, exploring alternatives and incorporating the results into subsequent inquiry. Feedback may prompt revision both of hypotheses and of the criteria used to select, compare, and deepen them. An objection may change not only the answer but also the procedure used to look for the next one.

The tool itself was also revised along the way. Part of its code was rewritten by the author based on architectural suggestions produced by the LLM during the investigation. The model proposed changes; the author examined the proposals and decided how to incorporate them into the implementation.

Human × AI symbiosis.

The architecture uses two complementary criteria for selecting and reviewing solutions:

1. Autonomous flow: the initial directive seeks the most coherent, simple, and comprehensive explanation, with the least possible dependence on exceptions, compatible with the available evidence and resistant to the objections examined. On identifying contradictions or gaps, the system can retrace its steps and revise the evaluation weights and investigative criteria applied to the problem's boundaries. The process confronts not only the hypotheses but also the methods that produced them and the interventions of human dialectic, until the cycle ends with the best alternative found under the adopted criteria and limits for the objective, within the examined conditions.

2. Human dialectic: at the end of each cycle, the solution found is submitted to Socratic debate with the expert operator — a scientist or analyst. The objections and perspectives offered are not accepted automatically: when they provide plausible grounds for further inquiry, they are incorporated into the next cycle and subjected to the same criteria, with the aim of mitigating the tendency toward sycophancy. A clear human statement of the ultimate objective is as important as the dialectic itself.

LLM model capabilities × AI logical structure

At the beginning of the thesis's development, the author's architectural core comprised more than 270 instructional functions, distributed across approximately 410,000 characters of instructions and organized into modules, criteria, cycles, and review controls. These numbers indicate the scale of the structure used; they do not constitute evidence of quality or effectiveness. The method coordinated these procedures into investigative cycles rather than depending on a single prompt. These functions were conceived to reduce known limitations of LLMs on open-ended tasks, such as typicality bias, response homogenization, and diversity collapse (or mode collapse). These tendencies lead models to favor familiar and semantically similar responses — often recombining existing patterns or, to please the operator, generating answers and solutions without evidential support — instead of exploring genuinely divergent and novel alternatives.

The system functions as an epistemic scaffold: it organizes exploration, comparison, and revision carried out with existing LLMs. The methodological hypothesis is that this organization may reduce dependence on the nominal capabilities of the base model. Investigating whether intermediate models structured by this method can outperform more advanced models used without the same architecture requires comparative evaluations. The proposal is to examine how much the quality of inquiry also depends on how the process is organized.

This tool conducted the inquiry presented in the thesis. In Chapters 9 and 10, the TMC formulates an architectural proposal inspired by the organization of living cognition under constraints: investigating how the effects of one operation may modify selection, retrieval, evaluation, and effort in subsequent operations. The tool used and the proposed Apparatus have distinct roles; using the former does not demonstrate implementation of the latter.

Topological method of the AI algorithm.

The method operates through four recurring stages, arranged in iterative cycles, with returns to earlier stages and exploration of alternatives:

a) Suspend: provisionally bracket inherited meanings associated with the object under investigation, preventing earlier definitions from predetermining the conclusion.

b) Reframe: replace the isolated object with the process, Trajectory, or system that produced it and continues to condition it. Reframing varies the scale, instance,

system boundary, and temporal window of analysis, constructing and comparing alternative hypotheses.

c) Integrate: confront the reframed hypotheses with knowledge drawn from different fields, seeking convergences, complementarities, incompatibilities, and contradictions.

d) Stress-Test: confront the proposal through two complementary dialectics. In simulated dialectic, the system subjects hypotheses to rival theories, objections, borderline cases, and alternative propositions. In real dialectic, the system and human operator confront interpretations, criteria, interim findings, and the very objective of the inquiry.

Authorship and the use of artificial intelligence tools

The authorial organization of modules, criteria, and cycles was developed under the author's direction, with assistance from LLMs in exploring and revising alternatives, including architectural suggestions. Decisions on what to incorporate and responsibility for the work remained human.

The process of investigation and theoretical structuring was assisted by language models operating under strict human methodological oversight. In this work, the author's protocol guided the exploration of alternatives, confrontation of objections, and revision of formulations. The process allows the contribution of this organization to inquiry to be examined, without establishing, by itself, superiority over other methods or models. Artificial intelligence tools provided instrumental support in organizing, comparing, reviewing, and exploring alternatives, without replacing authorship, conceptual decision-making, source verification, or the author's argumentative responsibility.

The theory's entire authorship, including its premises, logical structure, and conclusions, belongs exclusively to the human author, who assumes full epistemological, theoretical, and ethical responsibility for the thesis's content. The use of these protocols is presented as a possible aid to frontier research, whose contribution to rigor and creativity must be assessed across different problems, models, and conditions of use.

1.4 — MAP

How the thesis is organized

1.4.1 — STRUCTURE OF THE THESIS

READING MAP · FIVE STAGES OF THE SITE

THE PROBLEM OF CONSCIOUSNESS `01 · 02a · 02b · 02c`

Why is there felt experience, and why does explaining how a system functions not settle that question?

THE MINIMUM OF CONSCIOUSNESS `03 · 04 · 05 · 06`

The Triad examines three kinds of work involved in living continuity: maintaining oneself (Inside), being transformed through encounters (Between), and continuing beyond the instance (Ahead).

THE UNITY OF CONSCIOUSNESS `07 · 08`

How an organization constitutes its unity, recognizes aspects of itself, and articulates thought, emotion, and feeling.

AI AND CONSCIOUSNESS `09 · 10`

What the Trajectory of living beings may teach us about building an AI capable of transforming the way it learns.

CONSCIOUS LIFE `11`

What defines a living system when its matter, form, and rhythm can change?

The groupings guide navigation; they do not mean that life, cognition, and consciousness are equivalent.

The thesis is organized into four cumulative movements:

1) DELIMIT — Chapters 1–3.

Introduces the problem, destabilizes the presumed separation between function and feeling, returns to life, and formulates the Triad.

2) DECOMPOSE — Chapters 4–6.

Didactically separates Inside, Between, and Ahead to show the causal work of each direction without treating them as independent modules.

3) INTEGRATE — Chapters 7–8.

Brings the Triad together in Ente Eu, extends it to Nós and Linhagem, and investigates abstraction, efficiency, thinking, and feeling under constraint.

4) TRANSLATE AND TEST — Chapters 9–11.

Confronts AGI architectures, formalizes the Apparatus, and uses the Triad as a measure for artificial life, borderline cases, and conditions for revision.

5) STRESS-TEST — PHILOSOPHICAL CONFRONTATION

During the TMC's development, debates simulated on the basis of the positions of David Chalmers, Daniel Dennett, Frank Jackson, Ned Block, Thomas Nagel, and Giulio Tononi contributed to the investigation and revision of the hypotheses.

These parallel debates produced conceptual convergences, reframings, and counterpoints among different theories. Their purpose was not to adopt a particular philosophical school, but to confront perspectives and examine what each required preserving, conceding, or revising.

The original records were preserved and may be presented in the future as historical documentation of the inquiry and of changes in the AI architecture.

The thesis presents, in Chapter 2, the selected debate based on Chalmers's positions. These exchanges are simulations; they do not imply direct interaction with the authors or validation of the TMC by their positions.
