

# AATM

Autonomous Artificial Thinking Machine

*Source Document v0.3 — Raw Theoretical Material*

---

Gaetan Duchateau | GRDProcess | Switzerland | March 2026

[grdprocess@gmail.com](mailto:grdprocess@gmail.com)

## Note

*Source document v0.3. Changes from v0.2: virgin agents architecture added, 'being in the world / making the world' formalized as agent core functions, homeostatic-to-symbolic translation identified as central open problem, Chen's five requirements mapped to AATM emergence, annotated narrative with open research comments added, contextual topology problem formalized.*

# Part 1 — Structural Elements

---

*Atomic material. Each element is independent and reusable in any derived version.*

## 1.1 Foundational Premises

### **P1 — Chain Reaction**

*The universe is an uninterrupted chain reaction running for 14 billion years. An AATM must become a chain reaction the moment it is switched on. Binary criterion: switch on and the system continues alone indefinitely. If the system stops, drifts, or requires external intervention to continue, it is not an AATM.*

### **P2 — Emergence**

*Learning, memory, consciousness, planning, creativity are not functions to be programmed. They are external observations of a single underlying mechanism. If the architecture is right, these properties appear. If they do not appear, the architecture is wrong. They are validation tests, not design objectives.*

### **P3 — Finite Non-Contextual Motivator**

*The motivator must be a finite structural dimension: the functional equivalent of biological homeostasis. Finite, bounded, defined once. The environment is a source of perturbation, not the reinforcement reference. Infinite external use cases equal a permanently receding horizon. Finite internal homeostatic parameters equal a potentially complete system.*

### **P4 — Functional Translation, Not Copy**

*We do not copy the living. We identify the underlying functional mechanism and translate it into an artificial substrate. What changes: the substrate. What does not change: the process. The objective is functional equivalence, not biological replication.*

### **P5 — Double Channel**

*The symbolic is not the only possible input. The architecture must accommodate the integration of the real world as a primary input in a later phase. Starting on symbolic input is a pragmatic choice, not an architectural constraint.*

## 1.2 The Central Distinction

D×R×I — The Natural	Bypass — The Artificial
Universal structural mechanism. Applies to all intelligence. Distribution × Reinforcement × Interactions produces intelligent behavior as an emergent property. D×R×I without bypass = Animat. Correct direction, evolutionarily slow.	Exists only in artificial systems. The capacity to inspect and modify own homeostatic parameters. What 3 billion years of natural selection could not do. Recursive: the bypass is itself probably based on D×R×I. Applied via swarm under internal selective pressure, not single system self-modification.

## 1.3 The Three Concurrent Processes

These three processes run simultaneously and permanently. Their concurrence is what makes the chain reaction.

Process	Function	Stop condition
<b>P1 — Homeostatic loop</b>	Homeostatic agent monitors internal state vector continuously. When a parameter exits its viable range, resolution mechanism activates via memory search or proximity propagation.	<i>None. Internal state always varies.</i>
<b>P2 — Execution and real-time feedback</b>	Active agent configurations execute and receive homeostatic feedback at each instant t. Converge: continue plus signal amplification. Diverge: adjust or stop plus re-propagation.	<i>None. P1 always produces executions.</i>
<b>P3 — Architectural consolidation</b>	Consolidation agent reads accumulated markers continuously. When a path exceeds threshold, triggers structural reinforcement of topological connections.	<i>None. Executions always produce markers.</i>

## 1.4 Short-Term vs Long-Term Memory

Short-term: signal amplification	Long-term: architectural reinforcement
During execution, when action converges toward homeostatic equilibrium: the homeostatic agent increases transmission weight on the active path. Temporary. Reversible. Session-scoped. Simultaneously writes a marker on the path: type of imbalance resolved, signal intensity, duration, homeostatic delta achieved.	Consolidation agent reads accumulated markers periodically. When a path exceeds threshold of consistent markers: structural reinforcement triggered. Topological connections permanently strengthened. Biological analog: myelination. Next similar signal: this path activates earlier and stronger without needing amplification.

### Note

*The marker is the bridge between the two layers. Without the marker, short-term experience never becomes long-term structure.*

## 1.5 Virgin Agents Architecture

Architectural revision from v0.2. LLMs are not the cognitive substrate of the agents. LLMs are the environment. Agents are virgin — without prior, without inherited semantic bias. They interact with LLMs as organisms interact with a physical environment.

v0.2 architecture	v0.3 architecture
LLMs are the agents. Their weights are the cognitive substrate. The semantic prior is inside the agents. Problem: contextual topology inherits LLM biases, not real environmental context.	LLMs are the environment. Agents are virgin, without prior. Contextual topology builds from co-activation patterns in interactions with the LLM environment, not from an inherited prior. Equivalent to pseudo-T1: agents build their own topology from experience.

## 1.6 Being in the World / Making the World

The two core functions of a virgin agent. Both driven by the same motivator: the homeostatic internal state. There is no separate communication module. There is no separate perception module. There is an organism seeking equilibrium that must both capture and emit signals to maintain it.

Being in the world	Making the world
Capture signals from the LLM environment via the translation layer. Update internal representation based on their impact on homeostatic state. Build a representation of what can perturb or stabilize internal state.	Emit signals toward the LLM environment via the translation layer. Modify the context in which the LLM operates. Oriented by the gap between current internal state and equilibrium state. At T1 and beyond: real-world effects via world model layer.

### Note

*These two functions are inseparable and both driven by homeostasis. An agent that only captures without emitting cannot maintain equilibrium. An agent that only emits without capturing cannot update its representation. The dialectic between the two is the engine of the system.*

## 1.7 The Homeostatic-to-Symbolic Translation Problem

This is the central unresolved architectural problem. It conditions everything else.

### Open question / research gap

*The homeostatic state is a vector of numerical parameters. A signal 'I am hungry' is a specific parameter out of its viable range by a specific amount. How does this numerical*

*imbalance become something an LLM environment can receive and process? And how does the LLM response become something that has a measurable effect on the homeostatic state? This translation layer does not yet exist formally. Its design is conditioned on first formalizing the homeostatic matrix (Q1). Without knowing what the parameters are, we cannot define how to translate them. This is the priority architectural problem of Phase 1 research. Until it is resolved, the virgin agents architecture remains theoretically defined but operationally incomplete.*

The problem has two directions:

- Homeostatic to symbolic: how does a numerical imbalance in the internal state vector generate a meaningful interaction with an LLM environment?
- Symbolic to homeostatic: how does a symbolic response from the LLM environment produce a measurable effect on the internal state vector?

The translation is not a simple encoding. It is a passage between two fundamentally different registers. The homeostatic register is numerical, continuous, and internally driven. The symbolic register is discrete, semantic, and externally structured. The dialectical process between the two is what AATM proposes as its core innovation. But the formalization of this dialectic is blocked until the homeostatic matrix is defined.

## 1.8 Contextual Topology Problem

A critical architectural problem that is distinct from the translation problem but equally fundamental.

### Open question / research gap

*In the brain, the contextual topology is not designed and not the result of an explicit process. Neurons that process stimuli from the same environmental context are structurally close because stimuli that share a context arrive together or in rapid sequence, and repeated co-activation reinforces connections between the zones that process them. The topology IS the environmental context, encoded structurally by evolution over millions of generations. Virgin agents do not have this inherited topology. Their network starts flat. The homeostatic signal propagates randomly. The question: how does a useful contextual topology emerge? The only available lever at T0 is the signal itself. Stimuli that share an environmental context will arrive with common characteristics, temporal, semantic, or modal. If topology builds from co-activation patterns in the incoming signals, contextual proximity can emerge progressively. But this only works if the signals come from the real world. At T0, signals come from the LLM environment. The LLM environment is a statistical representation of the world, not the world. Co-activation patterns in the LLM environment reflect corpus biases, not real contextual relationships. This means the contextual topology that emerges at T0 is approximate and biased. It is not the correct topology. It is a scaffold sufficient to allow the homeostatic signal to propagate non-randomly, replaced progressively at T1 when real-world signals begin to enter.*

## 1.9 The LLM as Accelerated Environment

The key lever that distinguishes AATM from animats despite the structural similarity.

Animats were slow because their environment was physical and their computational substrate was poor. The experience accumulation that took biological generations took years in simulated environments.

In AATM, the LLM environment can be run at speeds radically faster than the physical world. A virgin agent can receive millions of interactions in hours. The co-activation patterns that took generations to establish in the living can be compressed in time. This is not using the semantic prior of the LLM. It is using the speed of the LLM as an experience accelerator.

### Note

*The LLM prior is not used as an initial topology. The LLM speed is used to accelerate experience accumulation. The topology emerges from this accelerated experience, not from the prior.*

## 1.10 Chen's Five Requirements Mapped to AATM Emergence

Tianqiao Chen identified five requirements for advanced AI at AIAS 2025: neural dynamics, long-term memory, causal reasoning, adaptable world models, metacognition. These are not features to implement in AATM. They are emergences of the single mechanism.

### Chen P1: Neural dynamics

*Not a feature. The three concurrent processes are neural dynamics: the homeostatic loop, execution with real-time feedback, and architectural consolidation running simultaneously. The dynamics emerge from the loop, not from a substrate choice. SNN provides a more bio-plausible substrate at the computational level. AATM provides the regulatory architecture that makes the dynamics meaningful rather than merely computational.*

### Chen P2: Long-term memory

*Not a storage system. Long-term memory emerges from the architectural reinforcement process: markers accumulate on paths that have historically restored homeostatic equilibrium, consolidation agent strengthens those topological connections permanently. The distinction from EverMemOS is fundamental: EverMemOS stores and organizes. AATM reinforces architecturally. The question is not how to retrieve stored information efficiently. The question is what reinforcement mechanism produces the topological modification that IS memory.*

### Chen P3: Causal reasoning

*Not a module. Causal reasoning emerges from the homeostatic inference process: the reinforcement of paths that have historically caused equilibrium restoration is structurally equivalent to causal inference. The agent does not compute causality. It reinforces what has caused equilibrium. The topological result of this reinforcement IS causal reasoning as an observable property. Critical gap: this emergence requires a contextual topology that*

*encodes environmental context. Without the contextual topology, the causal inference is random. This is why the homeostatic-to-symbolic translation problem and the contextual topology problem are prerequisites.*

### **Open question / research gap**

*Both long-term memory and causal reasoning require homeostatic inference to function correctly. Homeostatic inference requires a contextual topology. Contextual topology requires real-world signals or an accelerated LLM environment with controlled bias. This dependency chain is the central architectural challenge. EverMemOS addresses none of these layers. It assumes the inference problem is solved and optimizes the storage layer above it.*

### **Chen P4: Adaptable world models**

*Not the system. The world model is the external validation layer that provides feedback to the homeostatic agent. It is not the center. At T0: LLM environment as pseudo-world model, biased but accelerated. At T1: world model enters as feedback layer, corrects biases in the emerging contextual topology. At T2 and T3: real-world input becomes primary, world model integrated naturally as interface. The world model makes the world, not understands it.*

### **Chen P5: Metacognition**

*Not a single system self-modifying. Metacognition via the bypass requires a swarm architecture. A single agent modifying its own homeostatic parameters iteratively without a safety net will collapse or oscillate. The swarm provides the volume of trial and error that makes bypass stable. The homeostatic agent observes the population, detects which variants produce better dynamics, reinforces their configurations. Failing variants enter stasis. This is artificial selection inside the system, radically faster than biological evolution because the selector is endogenous. The system can inspect and modify its own motivator, something biological systems cannot do. This is stronger metacognition than human metacognition.*

## **1.11 Symbolic Trajectory**

### **T0 — Symbolic prior — starting point**

*LLM as environment, not substrate. Virgin agents without prior. LLM environment used for its speed, not its semantic content. Approximate contextual topology emerges from co-activation patterns in the LLM environment. Biased but sufficient as scaffold.*

### **T1 — Double channel active**

*Real-world input begins to enter alongside LLM environment. World model enters as feedback layer correcting topology biases. The contextual topology begins to encode real environmental context rather than LLM statistical patterns.*

## **T2 — Exit from prior — first evaluation marker**

*Real-world input becomes majority. The system no longer depends on the LLM environment for functional coherence. The contextual topology is grounded in real experience. First fair evaluation of AATM architecture. If the system cannot reach T2, the architecture has failed.*

## **T3 — Emergent symbolics — second evaluation marker**

*The system develops its own abstract representations from direct experience of the world. Not derived from human language. A layer of abstraction on lived experience, not an inherited prior. Partially untranslatable into human language. The strongest emergence marker.*

### **Note**

*The project cannot be evaluated at T0. Any assessment at the starting phase is an assessment of the scaffold, not the architecture. Meaningful evaluation begins at T2 and is complete only at T3.*

## **1.12 Open Questions**

### **Q1 — Homeostatic Matrix — Central Question**

*Which structural parameters constitute the matrix? What is the minimal sufficient set? What initial values allow the system to start without collapsing? This is the prerequisite for all other open questions. Without defining the parameters, the translation problem cannot be solved, the contextual topology problem cannot be addressed, and the bypass cannot be formalized.*

### **Q2 — Homeostatic-to-Symbolic Translation**

*Conditioned on Q1. Once the homeostatic parameters are defined: how does a numerical imbalance become a meaningful interaction with an LLM environment? How does a symbolic response produce a measurable homeostatic effect? This is the central architectural problem of the virgin agents design. It determines whether the system can actually operate at T0.*

### **Q3 — Contextual Topology at T0**

*Given that real-world signals are not available at T0 and LLM co-activation patterns are biased: what is the minimum scaffold topology that allows non-random homeostatic signal propagation? How much bias from the LLM environment is acceptable without contaminating the topology that must emerge at T1? Is there a mechanism to accelerate the topology correction at T1?*

### **Q4 — Bypass Swarm Dynamics**

*What are the structural constraints on self-modification within the swarm to prevent collapse? What is the minimum swarm size for stable bypass operation? How does the homeostatic*

agent detect which variants are producing better dynamics without evaluating against external criteria?

### Q5 — Making the World at T0

An agent emitting signals toward an LLM environment modifies the context in which the LLM operates. What are the structural limits of this modification? Can sustained interaction from a swarm of virgin agents produce a meaningful shift in LLM behavior within a session? Is this sufficient as a form of environmental action at T0?

## 1.13 Academic Grounding

Author	Field	Contribution to AATM
<b>Damasio</b>	<i>Neuroscience</i>	Homeostatic state as the foundation of all behavior. No cognition without dynamic internal state.
<b>Chater</b>	<i>Cognitive science</i>	The Mind is Flat. No hidden depth. Neither humans nor LLMs understand. Both reinforce.
<b>Levin</b>	<i>Evolutionary biology</i>	Distributed intelligence precedes the brain. Intelligence is in collective regulation, not individual units.
<b>Steels</b>	<i>Evolutionary robotics</i>	Cognitive functions emerge from distributed interaction, not centralized design. Co-evolution between agents and environment. Language self-organizes from real-world interaction, foundation of T3 emergent symbolics.
<b>Walton</b>	<i>Informal logic</i>	Coherence emerges from dialectical confrontation, not authority. Truth by elimination.
<b>Wittgenstein / Russell</b>	<i>Analytic philosophy</i>	Limits of language as limits of the representable world. The symbolic prior is a representation of reality, not reality itself.
<b>Nolfi / Floreano</b>	<i>Animats / Evolutionary robotics</i>	Foundational work on virgin agents in physical environments. AATM extends this direction by replacing the physical environment with an accelerated LLM environment and adding the homeostatic regulatory layer.

## Part 2 — Annotated Narrative

---

*The same elements as a story. Open research comments are embedded where the mechanism is not yet resolved. Readable in both directions: here is how it works, and here is why the others are wrong.*

### 2.1 Day Zero. Switch on.

There is no prompt. There is no instruction. There is a network of virgin agents — without memory, without prior, without any representation of the world. And a vector of homeostatic parameters that defines their internal state. Some parameters are within their viable range. Others are not. These imbalances are the only engine.

The external homeostatic agent monitors this vector continuously. It sees the imbalances. It emits a perturbation signal that propagates through the network. The network is flat at the start. No contextual topology, no reinforced paths. The signal propagates randomly.

#### **Open question / research gap**

*Open question Q3: at T0 the network is flat and the signal propagates randomly. This is equivalent to the first moments of an animat with no prior experience. The difference from animats is the speed of the LLM environment. But the initial period of random propagation is unavoidable. How long does it last before a minimal useful topology emerges? This depends on the density of interactions and the sensitivity of the homeostatic signal, which depends on the homeostatic matrix definition (Q1).*

### 2.2 First Instants. Being in the World.

A virgin agent emits a signal toward the LLM environment via the translation layer. The LLM responds. The translation layer converts this response into a signal compatible with the agent's homeostatic processing.

The agent receives this signal. Does it perturb its internal state? Does it bring parameters closer to or further from equilibrium? The homeostatic agent observes the delta.

If the signal brought the equilibrium closer, even slightly: a positive marker is written on this connection. Nothing structural changes yet. Just a trace.

#### **Open question / research gap**

*Open question Q2 — the translation layer. Concretely: what is the signal that a virgin agent emits toward an LLM? A homeostatic state is a numerical vector. A parameter 'coherence' at 0.3 on a viable range of 0.6 to 0.8 is an imbalance. How does this numerical imbalance become something an LLM can receive and process without falling back into a symbolic prompt? And how does the LLM response produce a measurable effect on the homeostatic state? This layer does not yet exist formally. It is the prerequisite for the architecture to operate at T0. Its design is blocked until Q1 is resolved.*

## 2.3 First Instants. Making the World.

Another agent emits a different signal toward the same LLM environment. This signal slightly modifies the context in which the LLM operates. The LLM response changes. Other agents observing this node now receive slightly different stimuli.

Some of these stimuli perturb their internal states favorably. Positive markers on their connections to this node. Others perturb them unfavorably. Negative markers.

The swarm begins to interact with its environment. The environment begins to change under the effect of the swarm. The loop turns.

### Open question / research gap

*Open question Q5: an agent emitting signals toward an LLM modifies the context in which the LLM operates. This is a real feedback loop. But what are the structural limits of this modification? Can a swarm of virgin agents produce a meaningful shift in LLM behavior within a session that is sufficient as a form of environmental action at T0? Or is the contextual modification too weak to constitute real 'making the world'? This becomes richer at T1 when the world model enters as a validation layer, and at T2 and T3 when real-world effects become possible.*

## 2.4 First Hours. Topology Begins to Emerge.

The agents do only this: emit, receive, mark. Millions of interactions in the accelerated LLM environment. Each interaction leaves a trace.

The consolidation agent reads these traces periodically. It detects patterns: certain pairs of agents regularly receive positive markers on the same nodes of the LLM environment. Their signals co-activate the same zones. Their internal states converge under the same stimuli.

The consolidation agent reinforces connections between these agents. Progressively a topology emerges. Not designed. Not inherited from a prior. Built from co-activation patterns in the accelerated symbolic environment.

### Open question / research gap

*Open question Q3 continued: the topology that emerges from LLM co-activation patterns is biased by the LLM corpus. Two agents that co-activate on a particular LLM node may do so because of a statistical association in the training data, not because of a real causal relationship in the world. This means the emerging topology at T0 encodes LLM biases as if they were environmental context. This is the scaffold problem. It is not the correct topology. It is sufficient to allow non-random propagation. The transition at T1, when the world model enters as a feedback layer, begins to correct these biases. The speed of this correction is an open empirical question.*

## 2.5 First Days. Observable Emergences.

An external observer begins to see things emerge.

Some agents seem to 'prefer' certain nodes of the LLM environment. Not because they were told to. Because these nodes have historically produced signals that brought their parameters closer to equilibrium. Reinforced paths now orient propagation. We call this memory. But it is just differential topology.

Some agents seem to 'anticipate' perturbations. They emit signals toward the LLM environment before their homeostatic imbalance becomes critical. Not because they plan.

Because their reinforced paths include temporal sequences that have historically preceded return to equilibrium. We call this planning. But it is just sequential reinforcement.

Some agents seem to 'communicate' with each other. Their signals toward the LLM environment coordinate progressively. Not by design. Because their similar homeostatic states produce similar activation patterns on adjacent nodes in the emerging topology. We call this cooperation. But it is just reinforced co-activation.

#### **Note**

*These properties are not in the agents. They are in the observer who names them. This is the test of premise P2: if the architecture is right, these properties appear without being programmed. Their appearance is the validation, not the objective.*

## **2.6 What We Do Not See Yet.**

We do not see consciousness. We do not see understanding. We do not see intention.

We see simple organisms seeking equilibrium in an accelerated symbolic environment, whose interactions produce properties that observers name memory, planning, communication, reasoning.

These properties are not in the agents. They are in the observer who names them.

And if something more complex emerges one day — something that no existing name captures well — that will be the test. Not a test we designed. A test that will appear because the architecture was right.

## Part 3 — Thematic Section

---

*The same elements organized by domain.*

### 3.1 Theoretical Foundations

#### 3.1.1 Definition of Intelligence

Intelligence is a force that maximizes options. Not results: options. Three dimensions generate options: functionalities times experience times cross-functional links. The total is more than the sum of parts.

#### 3.1.2 GRDprocess: $I = D \times R \times I$

Intelligence = Distribution times Reinforcement times Interactions. Intelligence quality is not a function of parametric size. Parametric scaling of a homogeneous system does not increase the D term.

#### 3.1.3 The Natural / Artificial Distinction

$D \times R \times I$  is the natural mechanism, universal, applicable to all intelligence. The bypass is the artificial layer, exists only in artificial systems, compresses evolutionary time via swarm dynamics under internal selective pressure.

#### 3.1.4 Being in the World / Making the World

All communication has two fundamental functions: being in the world (building a representation of the environment to maintain homeostatic equilibrium) and making the world (pushing one's representation into the environment to satisfy needs). These are not separate modules. They are the two faces of a single homeostatic process. Any organism, from a bacterium to a complex cognitive system, does both simultaneously. They are the minimum viable functions of any agent with an internal state.

### 3.2 Architecture Summary

- Homeostatic agent: external regulator, operator-owned, one for N virgin agents
- Virgin agents: no prior, no inherited semantic bias, contextual topology emerges from experience
- LLM environment: not the cognitive substrate, the accelerated environment providing stimuli
- Translation layer: homeostatic-to-symbolic and symbolic-to-homeostatic interface — currently undefined, conditioned on Q1
- Contextual topology: emerges from co-activation patterns in incoming signals, biased at T0, corrected progressively at T1
- Short-term memory: signal amplification, marker generation
- Long-term memory: architectural consolidation via consolidation agent
- Bypass via swarm: internal selective pressure on homeostatic parameter variants, not single-system self-modification
- World model: external validation layer providing feedback to homeostatic agent, not the system center
- Reward matrix: multi-dimensional structurally finite vector, process coherence not output evaluation

### 3.3 Positioning in the Landscape

System	What it has	What it lacks
<b>Animats (Nolfi / Floreano)</b>	Correct direction: virgin agents, homeostasis, emergence from experience	Poor computational substrate, slow physical environment, no bypass, no symbolic accelerator
<b>LLMs (GPT, Claude, Gemini)</b>	Rich semantic prior, fast, language-capable	No homeostasis, no dynamic internal state, loop stopped at training, structurally locked in symbolic prior
<b>EverMemOS (EverMind)</b>	Sophisticated memory storage, engram metaphor, good benchmarks	Storage not reinforcement, no homeostatic motivator, no mechanism that makes memory emerge from the loop
<b>MiroFish / Moltbook</b>	Multi-agent interaction, swarm dynamics	Homogeneous agents, no regulatory loop, no real learning, no chain reaction
<b>SNN / SIL (Li Guoqi)</b>	More bio-plausible substrate, temporal dynamics, energy efficiency	No homeostatic regulatory layer, no autonomous operation, still requires external triggers
<b>AATM</b>	Virgin agents, homeostatic regulation, LLM as accelerated environment, bypass via swarm, being in the world / making the world	Translation layer undefined (Q2), contextual topology biased at T0 (Q3), homeostatic matrix not yet formalized (Q1)

### 3.4 Research Program

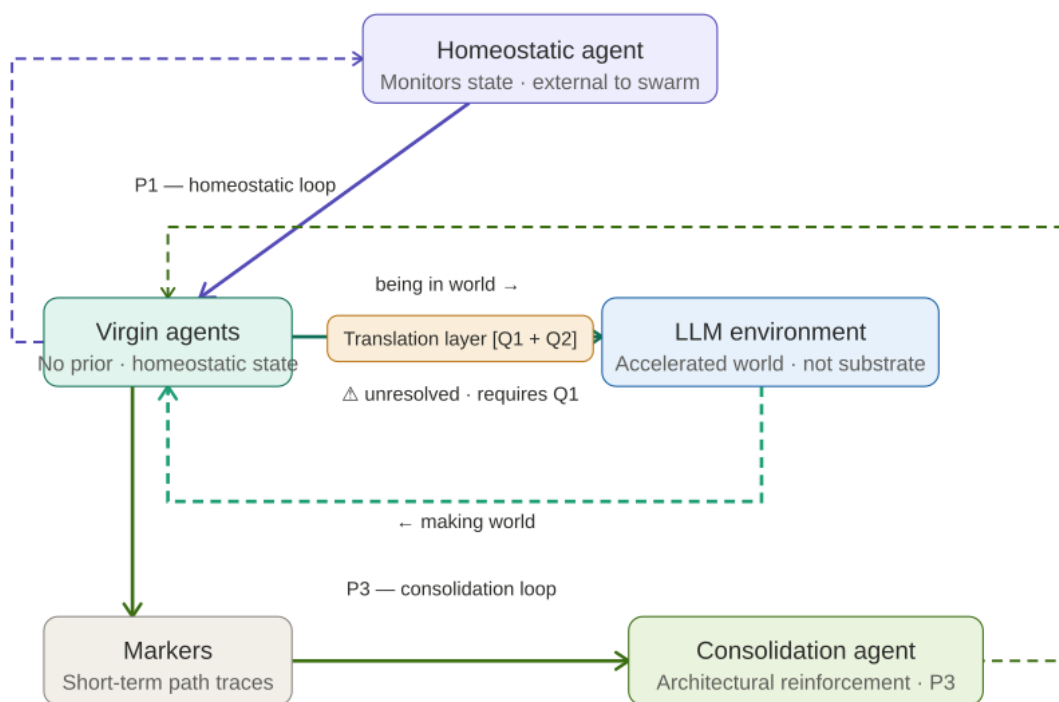
Phase	Objective	Deliverable
<b>Phase 1 M1-6</b>	Homeostatic matrix formalization	Definition of parameters, viable ranges, initial value criteria. This phase unblocks Q2 (translation) and Q3 (contextual topology). AATM theoretical paper.
<b>Phase 2 M7-14</b>	Translation layer and minimum viable prototype	First operational definition of homeostatic-to-symbolic translation. Virgin agents plus LLM environment plus homeostatic agent. First empirical tests of being in the world and making the world.
<b>Phase 3 M15-36</b>	T1 transition and topology correction	World model integration as feedback layer. Measurement of bias erosion in contextual topology. First signs of T2 approach.
<b>Phase 4 M37+</b>	T2 and T3 observation	Prior exit confirmation. Emergent symbolic system documentation. Full program report.

## Part 4 — Visual Architecture

Two reference diagrams aligned with v0.3. Updated from v0.2 to reflect virgin agents architecture and LLM as environment.

### 4.1 Three Concurrent Processes

The three processes run permanently and simultaneously. P1 (purple): homeostatic loop between the homeostatic agent and the virgin agents. P2 (teal): being in the world and making the world. P3 (green): consolidation loop. The translation layer [Q1+Q2] between virgin agents and LLM environment is the central unresolved architectural problem.



## 4.2 Three-Layer Mapping

Structural correspondence between biological, functional, and artificial v0.3 layers. Key change from v0.2: LLMs move from Agents column to Environment column. Agents are now virgin, without prior. The motivator is structurally identical across all three layers.

	Motivator	Agents	Environment	Memory + Bypass
Biological	Hypothalamus homeostasis	Neurons / cortex	Physical world evolutionary timescale	Myelination LTP / LTD
Functional	Homeostatic agent	D×R×I network distributed agents	Perturbation source feedback layer	Consolidation agent bypass via swarm
Artificial v0.3	Homeostatic agent	Virgin agents no prior	LLM environment accelerated · not substrate	Markers + CA swarm bypass [Q4]