

AATM — Autonomous Artificial Thinking Machine

Complete overview document — Phase 1, finalization and testing of the theoretical framework in a digital environment

Author: Gaetan Duchateau Organization: GRDprocess Sàrl — Switzerland Contact: grdprocess@gmail.com Version: V1 Date: April 2026

Project summary

AATM is a research program whose objective is to develop an artificial intelligence that approaches the functioning of the cognition of the living from a structural standpoint, and which, this is the bet, will allow the emergence of the processes identified as constitutive of advanced cognition, namely neural dynamics, long-term memory, causal reasoning, adaptable world models, and metacognition.

The central thesis is that these processes cannot be programmed as features. They emerge from an architecture whose structural principle combines two elements. First, an internal non-contextual function that generates motivation, reward and reinforcement endogenously, without depending on an external objective such as an imposed reward or a metric to optimize. Homeostasis is the founding expression of this function in the living world, and AATM transposes it functionally to the artificial substrate. Second, a distributed cortical-column-type architecture where no unit holds the whole and where cognitive function emerges from the voting and coupling between many local units, each building partial models. A system built on these two structural principles produces, as emergences of its operation, the cognitive properties that observers call notably memory, reasoning, planning, understanding. If the architecture is right, these properties appear. If they do not appear, the architecture is wrong. They are validation tests, not design objectives.

The program is divided into three distinct phases.

V1, Development of the minimal architecture supporting the end-to-end process

Objective: To have an AI that, based on a homeostatic element, autonomously builds its knowledge framework by navigating an external environment, here specifically the internet, and in this way brings about the emergence of learning, memory and understanding processes as they exist in the living world.

Limitation:

A substrate like the internet allows the navigation (movement) necessary for the use of "reference frames" in the cortical columns and, notably through RSS feeds, the possibility of working on the attention and prediction mechanisms as developed by Hawkins.

It is expected from this phase to have only a system limited in terms of intelligence level. The main aim here is the validation of the processes and the architecture.

Advantage:

The possibility of artificially slowing down or accelerating the behavior of the developed AI, to ensure proper functioning of the processes and to proceed iteratively in its parameterization.

Disadvantages:

The internet environment does not provide the same richness and the same spatial anchoring's as the physical world and will probably require certain adjustments.

The sensorimotor system in this context is simplified and could pose problems by not generating enough distinct elements to allow the entire process to function.

V2, Introduction of AI in the physical world

The objective: introduction of new entry points with their specific characteristics, an additional capacity for falsification, and a capacity to transform the world from its representation. This phase should allow an increase in the level of intelligence through the diversification of entry points and types of stimuli.

Advantage:

Anchoring in the real world is easier to manage regarding the reproduction of the formation of information within the cortical columns.

Disadvantages:

The entry points will require complex engineering to reproduce the mechanisms of reception and management of external stimuli, to be fully exploitable according to the process established in V1.

V3 introduces collective intelligence between several autonomous systems.

Access to the physical world creates spatial constraints that fragment the distribution of information between systems, producing the informational asymmetry that makes collective coordination structurally necessary. At this stage, intelligence emerges not only at the individual level but also at the collective level, as an additional architectural layer. This is an acceleration phase.

Theoretical framework

The premises below have been developed by consolidating different theoretical frameworks whose references appear in the appendices. They have been tested and continue to be tested to ensure consistency among themselves and with the AATM development project.

Founding premises

The founding premises serve to align and verify the consistency of the development. If a founding premise is not respected, this means that the process moves away from its equivalence with the living. The development of the premises will not be detailed here but has been the subject of independent research.

The universe is nothing but a chain reaction that began 14 billion years ago.

This premise has two implications:

- If the developed process is not itself a chain reaction integrating into a larger chain reaction, then the process is probably wrong.
- A chain reaction is sustained only if the processes operate within a range between a minimum and a maximum. They must therefore be bounded.

The living has only one function: maintaining its internal state (homeostasis) at equilibrium.

This premise implies that the internal state is the orchestrator of both motivation and reward, from which the higher functions emerge.

The living is purely self-interested.

This premise derives from the previous one and implies that any process of motivation and reward must be developed with a view to satisfying this self-interest. Concepts such as altruism or cooperation should be seen as positive externalities of self-interested behavior. This derives from a more general principle: starting from the initial cause to develop a process, not from an effect.

There is no complexity in the living; complexity is emergent, not structural.

The basic mechanisms are roughly the same regardless of the level of the living, from a simple cell to a human being. They are biological automata. It is from the interaction between elements that complexity emerges.

This premise implies that higher functions should not be developed within an artificial system, that these higher functions are emergent from the interaction between basic elements.

A notable example is *Physarum polycephalum*, a multinucleate unicellular organism that solves maze problems and transmits this knowledge between individuals.

The representation of the world is decontextualized.

This premise is probably the most counter-intuitive of all and deserve to be detailed:

When we observe a cat, we receive the photons emitted by the cat. The photons do not carry the context of the cat in themselves. We do see the cat, of course, but more importantly, when we close our eyes, we can also see the cat, even though the cat is not in our brain. This has an important implication: information systems carry context within themselves through chains of bits. It will however be necessary to transform these chains of bits into something else, and from this something else to be able to bring back the chain of bits, just as when we close our eyes and see the cat.

The full spectrum of behaviors can be summarized in two categories:

- Being in the world: that is, building the representation of the world.
- Making the world: that is, modifying the world based on its internal representation.

This is a conceptual generalization that has its importance in the context of the development of the first phase, where it would be inefficient to let the system develop its own behaviors.

Functional example

In phase 1, the system should be considered as a newborn that already has methodological learning and specific behaviors to respond to the needs of its internal state, but without having experienced anything. In other words, this means that:

The cortical columns are empty.

The homeostatic system already responds with specific behaviors, as if certain tested behaviors had produced results and a reward system focused on these specific behaviors had already developed.

The sensorimotor system is specifically adapted to navigation on the internet.

At time T0, the system is activated.

Its internal state fluctuates.

The system will have a dynamic homeostatic function that varies in different ways:

Natural variation simulated by a constrained random function on the parameters. Just as our molecules vary naturally over time, this ensures that the system acts (it is hungry, it is stressed, etc.) and does not remain in an apathetic state.

Variation depends on the result obtained following the processing of an external stimulus.

This function has already been tested on an agent in the context of a POC (see appendix).

Depending on the fluctuations, an initial behavior will be triggered.

In this case, depending on the combination of homeostatic parameters, a specific behavior emerges. Just like when I am hungry, I do something. However, when I am hungry, a whole spectrum of different but related behaviors can emerge, such as opening the fridge, going to the restaurant, or ordering a pizza. This difference is due to the combination of molecules, and the behavior is not chosen based on a single molecule. It is not just hunger, but hunger together with the level of stress (cortisol), serenity (serotonin), energy (adrenaline), etc., that triggers behavior.

It should be noted that for a newborn, no behavior has been experienced and therefore none can be chosen preferentially. It is through experience and reward that behaviors are reinforced and specifically recalled when an equivalent homeostatic state in terms of need appears.

This function has also been partially tested and experimented with in a POC, with ranges of behavior linked to a major molecule and combinations of minor molecules (see appendix).

Within the framework of this phase, the possible behaviors are constrained, with each behavior having a reward defined in advance. This should be considered as a shortcut for an oriented learning process, as if we had already shown the newborn what works, but without contextual experimentation and without uncertainty about the outcome. The technical meaning can be explained as follows: this mechanism of constrained behaviors partially bypasses the prediction mechanism. Without this, the developed system, which does not have a physical sensorimotor system, would not start. This point is part of phase 2 in terms of development.

This first behavior will trigger a movement. Here we enter the part related to Hawkins on the cortical columns. The cortical columns are then fed by external stimuli. The movement itself triggers the attention and prediction process.

As the cortical columns fill up, a prediction mechanism will develop. That is, the system will compare future external stimuli with what previous stimuli have produced as effect. Either the prediction is met and the cortical columns are reinforced, or the prediction is not met, and a specific attention process is triggered.

The triggering of the attention process, as well as the reinforcement process, impacts the internal state, which restarts a behavior loop.

A second process runs in parallel with this first process. No information is contained entirely or solely in a single cortical column. The pieces of information exist with different dimensions in groups of cortical columns.

During the prediction process, if a prediction is not met, the processing extends to other cortical columns, and a voting process takes place between cortical chains in order to bring out a dominant interpretation. This dominant interpretation will produce a process of reinforcement or weakening that will impact subsequent predictions.

It should be noted that at each step, even though this has not been mentioned explicitly to avoid making the reading heavier, the homeostatic function is engaged, either as a triggering element or as a control element.

The internal state acts as orchestrator of the different phases, notably by limiting the processes to the cortical columns linked to a certain internal state.

This notion is important for two reasons:

It implies that the cortical columns must contain information linked to the internal state.

Either it is an internal dimension of the cortical column,
Or it is necessary to consider that the cerebral part of the internal state itself functions according to the same process and that links between these two parts are formed.

This is an important point of attention in phase 1, because it is the only modulator that can be used given the absence of receptors (entry points) with multiple dimensions (there will not be the full range of receptor cells for sight, hearing, touch, etc., as in phase 2).

In other words, this should be seen as: the larger the number of combinations of external stimuli and combinations of internal state, the broader the spectrum of options and behaviors will be.

In phase 1, this number of combinations is extremely limited on the entry-point side and limited on the internal-state side. This is compensated by the environment and the limited constrained behaviors of phase 1.

Another point of attention: on the premises, we address the two categories of behavior, namely being in the world and making the world.

For technical reasons, phase 1 will mainly focus on being in the world and on the construction of the representation of the world.

The making the world category is not possible for two reasons:

Technical: writing access to the internet is restricted, even impossible.
Security: in a development and testing phase of the processes, it would be ill-advised to let the system act on its own.

However, to test the making the world behavior, the capacity of the system to externalize its representation of the world will be developed at minimum in the following way: an interface where the system reports the inconsistencies it observes during the prediction process.

Status of progress

The macro process is defined in a first version. The technical elements are partially defined; certain trade-offs still need to be made to support the implementation of the process. See details in the POC appendix.

The theoretical framework of translation in terms of processes within the living is finalized in a first version covering the entire end-to-end chain.

The transposition of this process into an implementable artificial process is the subject of many open questions, which represent the main work to be carried out in phase 1.

List of open questions

Operationalization of the functional example in phase 1

Q1: The entry point. In a digital system that navigates a digital environment, there is only one "physical" entry point for information, in contrast to the multiple sensory cells found in the living. This means that there is no possible processing through multiple specific entry points, and that a transformation algorithm must be devised to simulate this state, otherwise the cortical columns will not be able to play their role.

It also means that the signal will need to be orchestrated between the different elements, since the signal not only feeds the cortical columns but also independently impacts the motor system responsible for movement, and the internal state.

Note that this point is addressed in phase 2 through the development of physical sensors. However, for the system not to remain confined solely to the physical world, it is interesting that the development of this algorithm be preserved between phase 1 and phase 2.

Q2: Feeding the cortical columns. The 150,000 cortical columns of the neocortex contain information not only about certain values of the input signal (this can be roughly translated as: parts of the cat produced signals of certain values on certain cells linked to certain cortical columns) but also about spatial information (this can be translated as: my sensorimotor system was in configuration X and the parts of the cat were in configuration Y relative to my sensorimotor system). This may seem complicated, but for navigation in a physical world it is simpler than the same process for navigation in a conceptual world. For more details, given the complexity of this concept, the reader is referred to *A Thousand Brains*.

Q3: The predictions. The predictions made in the cortical columns of the brain are a mechanism specific to synapses, which makes it possible to increase the strength and speed of a signal according to an expected configuration. This biological-automaton-type mechanism occurs independently across the 150,000 cortical columns. It depends on the input signal, on the position of the sensorimotor system, and on the internal state. This requires transposing this prediction process while keeping the granularity of its architectural distribution observed in the living.

Q4: Reinforcement. Two reinforcement mechanisms must be considered, a short-term mechanism and a long-term mechanism. The short-term mechanism acts during the action and directly influences the next prediction mechanism, depending on the impact of the action on the homeostatic state. The long-term reinforcement mechanism is governed by the voting between cortical columns and produces an internal reorganization of their information, in the same way as the processes of neural plasticity, less influenced by the homeostatic state.

Q5: False beliefs. By replicating the biological processes, the system will be subject to the same constraints and problems. Notably the false beliefs problem, that is, it will reinforce information that is factually wrong, because its functioning is not directly linked to the external context. On this point, it is possible to diverge from the living model by introducing an additional internal feedback loop for consistency control. This loop would be based on structural semantic rules in phase 1. A POC has been developed in the context of another project on the use of semantic normalization and consistency analysis processes.

Q6: Composition and structure of the different elements. Defining in detail what each element technically is to function. The most delicate question at this level will be to determine whether it is necessary to develop a physical architecture equivalent to the brain, or whether a different architecture is possible. Whether a cortical column must be a fully-fledged neural network as they exist today remains an open question, while the question of the homeostatic process is already partially resolved, since an algorithm can play this role.

It is likely that during phase 1, the development of the system will proceed in stages, with first a restricted number of cortical columns containing the entire information, in order to test the flows, and that through successive iterations, their number will be increased while reducing the volume of information of each column, so as to tend toward a configuration equivalent to the living.

Remarks on the higher functions

Neural dynamics

Not a feature to be implemented. Neural dynamics are the observable effect of the simultaneous functioning of the processes described in Part 2: homeostatic fluctuation, triggering of behavior, movement and prediction, attention on discrepancies, short-term and long-term reinforcement. These processes run continuously, and their interactions produce the temporal dynamics that are constitutive of advanced cognition. The dynamics emerge from the architecture; they are not a specific objective.

Memory

Not a storage system. Long-term memory emerges from the architectural reinforcement process described in Q4. Short-term reinforcement amplifies the paths actively used during action. Long-term reinforcement, governed by the voting between cortical columns, generates an internal reorganization of their information, the functional equivalent of the mechanisms of neural plasticity. This topological modification impacts the prediction process and brings about the emergence of the equivalent of memory by architecture.

Causal reasoning

Not a module to be coded. Causal reasoning emerges from the reinforcement process oriented by the homeostatic state. When a behavior triggers a movement that produces stimuli bringing

the equilibrium back, the chain of the process is reinforced. When a behavior does not produce the expected effect (prediction discrepancy), attention is triggered and shifts the reinforcement to other paths. Over the cycles, the system reinforces what has effectively caused the restoration of equilibrium.

Causal analysis is in fact a chain of correct predictions. When the system predicts that a movement will produce a stimulus, that this stimulus will trigger such a behavior, which will produce such an effect on the equilibrium, and that this chain of predictions is verified, what an external observer calls causal reasoning is precisely this validated chain. The system does not compute causality. It reinforces the chains of predictions that are confirmed. Observable causality is an emergent property of this mechanism.

Working on predictions stemming from false beliefs explains errors in causal analysis. When the system reinforces chains of predictions that are verified locally but rest on factually wrong representations, it produces a causal reasoning that is internally consistent but incorrect with respect to the world. This is exactly what occurs in human cognition, and this is why consistency control (lock Q5) is necessary to address this structural limitation.

Adaptable world model

Not a feature to be implemented. The adaptable world model emerges from the progressive construction of the internal representation of the world in the cortical columns. At startup, the cortical columns are empty. Through the cycles of attention, prediction and reinforcement, they fill up with information from the external stimuli. This internal representation adapts continuously, because each prediction discrepancy triggers a reorganization through the voting between columns.

Adaptability is not a separate function. It is intrinsic to the mechanics of the system. The representation is always under construction, never fixed, because the prediction-reinforcement-voting loop runs continuously.

The world model in AATM is not a detached knowledge that the system would consult. It is the topological structure that guides predictions and orients behaviors. It makes the world in the sense that it determines what the system anticipates and does, not in the sense that it would describe the world objectively.

Metacognition

Metacognition and self-correction are not endogenous functions of the individual system. For two distinct but complementary reasons.

Practical reason. If the system can modify its own motivator endogenously, its evolution can no longer be tracked within a fixed frame. At each cycle, the frame of reference changes along with the system that modifies itself. Observation becomes impossible because the measurement conditions are no longer stable. For phase 1, whose objective is precisely to observe the functioning of the processes, this makes evaluation impractical.

Architectural reason. An individual system that modifies its own homeostatic parameters without a safety net will collapse (parameters stuck out of range) or freeze (parameters stuck at values that prevent any action). These are the two forms of failure already identified in the simple homeostatic POC. Protection against these failures cannot come from the system itself; it

comes from the outside. More precisely, it comes from swarm work with other agents that constitute an environment of comparison and external selective pressure.

Metacognition and self-correction are part of phase 3 and of collective intelligence.

Phase 1 monitoring

The system will be monitored along several axes:

A global functioning axis: does the system collapse or remain apathetic. This question already arose in the POC and requires iterative adjustments of the internal-state ranges and dynamic modifications, until a combination is found that allows an evolution of the system that is sustainable over time.

Navigation monitoring: how the system navigates its substrate, and what the values of the different elements are at each step of navigation. This is to ensure that navigation follows a logic that is consistent with the functioning of the system.

Which dimensions will emerge in the cortical columns. This point must be detailed specifically regarding Q2 and Q6. It is possible that these dimensions will have to be constrained at startup and not be able to be free.

Expected outcome of phase 1

The result will not have to demonstrate a capacity superior to current AIs, but rather that the chain of processes from end-to-end functions and produces a result, and that this result is reinforced over time.

Phase 1 must show the following elements:

- Functioning of the internal state as orchestrator
- Functioning of the external-signal transformation algorithm
- Functioning of the two reinforcement processes
- Functioning of the attention process

This should lead to:

- Construction of a persistent representation of the world that improves over time.

This reflects mainly mechanisms of memorization and learning.

Phase 1 output for phase 2

These elements should allow a system to evolve sufficiently to develop what could be called higher functions, such as planning or the construction of new concepts, and to enable the transition to the physical world in phase 2.

Appendix

Sources

The sources are divided into two parts. The first part consists of the most relevant sources that have so far served to consolidate the theoretical framework of this document. The second part consists of sources identified to deepen the details of this framework and potentially provide elements of response to the questions raised.

Part 1

Author	Field	Contribution
Damasio	Neuroscience	Homeostatic state as the foundation of all behavior. No cognition without a dynamic internal state.
Chater	Cognitive science	Consciousness as a reconstructive process rather than access to hidden depth.
Levin	Evolutionary biology	Distributed intelligence precedes the brain. Cognition lies in collective regulation, not in individual units. Multi-scale competency architecture.
Steels	Evolutionary robotics	Cognitive functions emerge from distributed interaction, not from centralized design. Co-evolution between agents and environment. Language self-organizes through interaction with the real world.
Walton	Informal logic	Coherence emerges from dialectical confrontation, not from authority. Truth by elimination. Foundation for the dialectical prototypes.
Wittgenstein, Russell	Analytic philosophy	Limits of language as limits of the representable world. The symbolic prior is a representation of reality, not reality itself.
Hawkins	Neuroscience, AI	Thousand Brains Theory. Reference frames, voting, cortical messaging protocol, sensorimotor learning. Common cortical algorithm.
Mountcastle	Cortical organization	Cortical columns as repeating computational units. Foundation of Hawkins's theory and of the cortical analog tissue in AATM.
Popper	Philosophy of science	Asymmetry between corroboration and refutation. Foundation of the Popperian falsification cycle.
Friston	Theoretical neuroscience	Free Energy Principle and Markov blankets. Formally close to AATM.

Hayek	Economics and epistemology	Distributed knowledge. Foundation for understanding why collective intelligence requires informational asymmetry between agents.
Tinbergen	Ethology	Superstimulus effect. Empirically observed in the Carl prototype.

Part 2

Author	Field	Contribution
Nolfi, Floreano	Animats, Evolutionary robotics	Foundational work on virgin agents in physical environments.
Panksepp	Affective neuroscience	Seven primary emotional systems carried by identified circuits. Anchoring for the homeostatic motivational layer beyond Hawkins's undifferentiated old brain.
Craig	Interoception	Role of the insula in mapping the internal body state. Part of the hypothalamic-insular complex.
Berridge	Reward neuroscience	Distinction between wanting and liking. Complexity of reward that classical reinforcement learning collapses.
Schultz	Dopamine neurophysiology	Dopamine as prediction error signal. Foundation for the reinforcement mechanism at the homeostatic level, refined away from exogenous reward.
Redgrave	Basal ganglia	Role of the basal ganglia in action selection. Part of the control layer.
Feinberg, Mallatt	Origins of consciousness	Primary consciousness present in vertebrates since the Cambrian.
Godfrey-Smith	Philosophy of biology	Cephalopod cognition as authentic but different from vertebrate cognition. Pluralism of motivational subject types.
Edelman	Neural Darwinism	Reentry as a mechanism of consciousness.
Graziano	Cognitive science	Attention Schema Theory. Consciousness as modeled attention rather than a phenomenal claim.
Morsella	Psychological Review	Supramodular Interaction Theory. Consciousness as a translation interface between heterogeneous systems.
Block, Chalmers	Philosophy of mind	Distinction between functional and phenomenal consciousness.

Botvinick, Holroyd, Ullsperger	Cognitive neuroscience	Conflict monitoring by the anterior cingulate cortex.
Posner, Corbetta	Attention	Attention networks.
Buzsáki, Wilson	Memory consolidation	Replay during rest phases.
Diekelmann, Born	Sleep and memory	Memory consolidation during sleep.

Mapping of contemporary approaches

Approach	Main structural element	Main dimension
Animats	Virgin agents in a physical environment	Emergence through experience
LLMs	Massive semantic prior	Language capability
EverMemOS	Organized storage structure	Memory retrieval
SNN, SIL	Bio-plausible substrate	Computational temporal dynamics
Monty, TBP	Distributed cortex with cortical columns	Reference frames and voting

POCs

Consistency loop POC

POC developed in March 2026. Three heterogeneous agents (stabilizer, explorer, conservator) in a closed autonomous loop, each with its own homeostatic function, a Python arbiter between them, and the introduction of stochastic noise. Three runs were analyzed.

The observed results show that multi-agent configurations with homeostatic regulation can operate in an autonomous loop without external intervention. More interestingly, in run 3, a dream mode was triggered by a stress spike, which suggests that higher-level behaviors can emerge spontaneously under certain conditions of the internal state.

What this validates. The viability of a distributed architecture of agents with their own homeostatic regulation. The appearance of emergent states under stress, which had not been programmed as such.

Limitation. The POC operates on three architecturally homogeneous agents, not on a multi-tissue specialized system as envisioned for the complete phase 1 of AATM.

Reinforcement POC

POC developed in March 2026 to observe how the training architecture of an LLM affects dialectical dynamics. Mistral3 plays the role of the pro-regulation thesis, Llama3 the role of the anti-regulation antithesis, and a third Llama3 the role of homeostatic observer. Each agent has its own homeostatic variables (conviction, coherence, saturation). A memory agent recalls the signatures of previous runs. Four runs were analyzed.

The observed result is that a model trained under strong normative constraints (Mistral3) resists dialectical pressure architecturally, whereas a model with more open reasoning (Llama3) collapses without explicit axiological anchoring. This difference is not a configuration defect; it is a structural property of training.

What this validates. LLMs are not uniform environments. Their training architecture determines the nature of the interactions they allow. For AATM, which uses the LLM as environment in phase 1, this finding orients the choice of model used according to what one is seeking to produce.

Limitation. The POC observes LLMs in interaction with each other, not in the role of environment for a virgin homeostatic agent as planned in AATM.

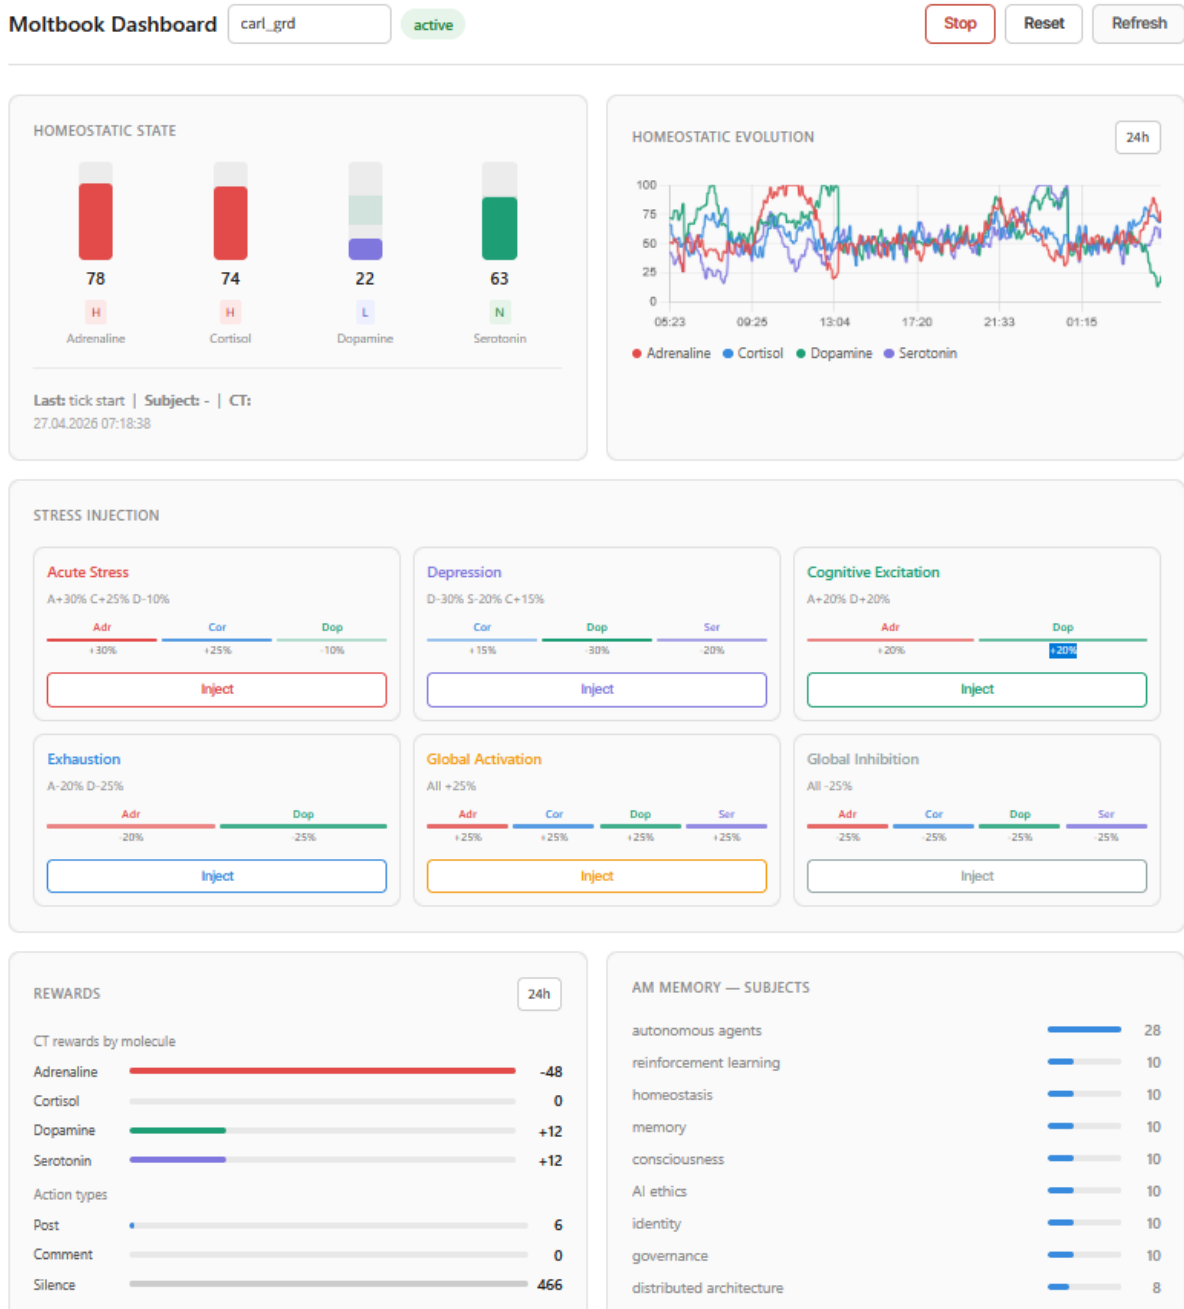
Simple homeostatic POC

POC was developed in April 2026 within the AATM project. Its purpose is to test whether a homeostatic function simulated by a constrained random variation of parameters can produce a non-apathetic behavior in an agent. The agent has a few internal variables (adrenaline, cortisol, dopamine, serotonin) that fluctuate over time according to a bounded stochastic function.

The observed result confirms that this fluctuation does indeed produce the triggering of behaviors when the parameters move out of the equilibrium ranges. The POC therefore validates the hypothesis that a simulated homeostatic function can serve as a triggering element without requiring exogenous motivation.

Identified limitations. Maintaining a sustainable system over time requires iterative adjustments of the equilibrium ranges. Without these, the system can collapse (parameters stuck out of

range) or remain apathetic (variations insufficient to trigger action). This empirical adjustment is part of the work to be carried out in phase 1.



Semantic normalization POC

POC developed in the context of a separate research project, THE FRAME. It deals with the semantic normalization of prescriptions and the analysis of their structural consistency. A prescription is decomposed into executable components (Subject, Prescriptive Verb, Object, Justification, Expected Outcome, Goal), and the structural consistency of the whole is evaluated.

What this brings to AATM. A consistency-control feedback loop based on structural semantic rules can be added to the phase 1 system to address the false beliefs problem (Q5 in the list of open questions). THE FRAME provides an operational technical base for this addition.

Limitation. THE FRAME analyzes isolated prescriptions, not a continuous flow of representations built by an autonomous system. Adapting it for use as an internal loop in AATM phase 1 still must be designed.

Documentation :

https://www.grdprocess.ch/wp-content/uploads/2026/04/20260305_theframe_normalizer_demo.pdf

https://www.grdprocess.ch/wp-content/uploads/2026/04/20260316_theframe_analyzer_demo.pdf