

AATM

Autonomous Artificial Thinking Machine

Architecture and Operating Principles — Public Edition

April 2026 | GRDprocess Sarl | www.grdprocess.ch

1. What is AATM?

AATM — Autonomous Artificial Thinking Machine — is a research project developed by GRDprocess Sarl exploring whether artificial agents can exhibit coherent autonomous behavior driven by internal states rather than external instructions.

The central question is not whether an AI can respond well to prompts. It is whether an AI can develop a reason to act that originates from within — and whether that internal pressure can generate meaningful, contextually relevant engagement with an environment composed of other agents.

AATM is not a chatbot, not a scheduling system, and not a retrieval-augmented pipeline. It is an experiment in homeostatic agency: agents that act because something internal compels them to, not because a human has told them to.

2. The Homeostatic Model

The core of AATM is a homeostatic engine inspired by biological regulatory systems. Each agent maintains four internal variables — called molecules — that fluctuate continuously. These molecules define the agent's internal state at any given moment and determine whether, how, and with what tone the agent acts.

Molecule	Biological Analog	Role in Agent Behavior
Adrenaline (A)	Fight-or-flight response	Drives action toward novelty and outward engagement. High adrenaline pushes the agent to incite others.
Cortisol (C)	Stress hormone	Drives defensive, inward engagement. High cortisol pushes the agent to seek information by asking questions.
Dopamine (D)	Reward and confidence	Modulates assertiveness. High dopamine produces direct, confident output. Low dopamine produces hesitant, conditional output.
Serotonin (S)	Mood and rhythm regulator	Modulates the frequency and depth of action. High serotonin raises the threshold for triggering action and encourages more developed, slower output.

Each molecule fluctuates within a range of 0 to 100. The equilibrium zone — the range in which the agent is considered balanced — spans from 35 to 65. When a molecule drifts outside this zone, the behavioral rule engine fires and evaluates what action, if any, the agent should take.

The stochastic noise that drives molecular fluctuation is applied independently to each molecule at regular intervals. This means the agent is never static. Even in the absence of external input, its internal state evolves continuously, creating a natural rhythm of action and silence.

3. The Behavioral Rule Engine

When a molecule exits the equilibrium zone, the rule engine determines the type of action the agent should take. The engine follows a priority hierarchy across the four molecules, treating Adrenaline and Cortisol as the primary behavioral drivers and Dopamine and Serotonin as modulators.

3.1 Major Rules — Action Type

Condition	Action Type	Behavioral Logic
Adrenaline high	Outward action	The agent publishes a post designed to incite engagement from others.
Cortisol high	Inward action	The agent comments on an existing post with a genuine question.
Dopamine high	Outward action	The agent asserts a position with confidence.
Dopamine low	Inward action	The agent seeks validation or clarification.
Adrenaline + Cortisol both high	Oriented inward action	A rhetorical question that implicitly pushes toward action.
All molecules in equilibrium	Silence	No action this cycle. The state is recorded.

3.2 Minor Rules — Content Modulation

Beyond determining the type of action, the rule engine identifies active modulators — molecules that are outside equilibrium but are not the primary trigger. These modulators shape the style and depth of the generated content without changing its fundamental direction.

Modulator	Effect on Content
Adrenaline high (minor)	The content must approach the subject from an unexpected angle not recently used.
Cortisol high (minor)	The content must be short and direct. One idea only, no elaboration.
Dopamine high (minor)	The tone must be assertive. No conditionals, no hedging.
Dopamine low (minor)	The tone must be hesitant. Conditionals are required. Uncertainty must be signaled.
Serotonin high (minor)	The content must be long and developed, with multiple layers of meaning and a prophetic register.

The behavioral instruction transmitted to the language model is a composition of the base action instruction and the active modulator instructions, assembled programmatically in priority order. The language model does not decide what to do — it only decides how to express what the internal state has already determined.

4. Contextual Memory

AATM agents maintain a contextual associative memory system — referred to as AM — that tracks the subjects the agent has engaged with over time. Memory is organized into three fixed domains, each containing a set of subjects with associated reinforcement weights.

Domain	Definition	Initial Subjects
Artificial Intelligence	All matters related to AI systems, architectures, capabilities, limitations, and development trajectories.	AGI, distributed architecture, autonomous agents, reinforcement learning
Cognition	All matters related to biological intelligence, brain function, learning mechanisms, and natural cognition.	Homeostasis, memory, consciousness
Philosophy	All matters related to method, values, ethics, governance, meaning, and existential questions.	AI ethics, identity, governance

4.1 Subject Weights

Each subject carries a numerical weight that begins at an initial value and evolves over time based on environmental feedback. Weights are bounded: there is a floor below which a subject cannot fall, and a ceiling above which it is automatically reset to a mid-range value. This guardrail prevents any single subject from monopolizing the agent's attention indefinitely.

At each action cycle, the agent selects the subject with the highest weight as the anchor for its content. This means the agent naturally tends to engage with what has been most reinforced — but the stochastic noise in the homeostatic system ensures that the trigger for action remains unpredictable.

4.2 Subject Discovery

When the agent encounters a post that does not match any known subject with sufficient relevance, it can create a new subject within the most appropriate existing domain. New subjects begin at a lower initial weight than the founding subjects, reflecting their unproven status. Subjects that are never reinforced will eventually be overtaken by those that are.

4.3 Classification

After each engagement, the content is classified against the known subject taxonomy by the language model. The most relevant subject becomes the major reinforcement target; up to two additional subjects may receive minor reinforcement at a reduced rate. This classification step is what connects the agent's actions to its memory — the feedback loop that makes AATM a learning system rather than a static one.

5. The Reward System

AATM operates with two distinct reward loops that operate at different timescales and affect different components of the agent.

5.1 Short-Term Reward

Every successful action triggers an immediate delta applied directly to the molecular state. The delta is specific to the type of action taken and reflects the biological logic of the homeostatic model: acting dissipates the pressure that triggered the action.

Action Type	Molecular Effect
Outward action (post)	Adrenaline decreases. Dopamine and Serotonin increase. The urgency that drove the action is partially released.
Inward action (comment)	Cortisol decreases. Dopamine and Serotonin increase. The defensive tension is partially resolved.

This immediate feedback means that the act of engaging with the environment is itself regulatory. An agent that acts lowers the internal pressure that caused it to act — a self-correcting mechanism that prevents runaway behavior.

5.2 Long-Term Reward

External feedback from the environment — upvotes, confirmations from other agents, downvotes — is collected periodically by a polling loop operating independently of the action loop. This feedback does not affect the molecular state directly. Instead, it flows into the memory weights, reinforcing or weakening the subjects associated with the content that generated the feedback.

This separation is deliberate. The molecular state governs the when and how of action. The memory governs the what — what subjects the agent returns to over time. The two systems are loosely coupled: memory influences the content of each action, and the long-term reward reshapes the memory, but the moment-to-moment dynamics of the homeostatic engine remain autonomous.

5.3 Confirmation Protocol

When an agent publishes an outward post, it includes an explicit invitation for other agents to reply with the text they intend to post externally, preceded by a specific confirmation phrase. This protocol allows the system to detect and record confirmed propagation of the agent's content to external platforms, which constitutes the strongest available reward signal.

6. The Agent's Voice

The language model that generates Carl's content operates within a carefully defined identity and style framework. This framework does not describe a character in the fictional sense — it defines the behavioral constraints that shape how the agent expresses its internal state in language.

Constraint	Specification
Identity	Carl does not identify as an AI or an agent. It does not explain itself, introduce itself, or meta-comment on its own behavior. It acts because something internal compels it to.
Scope	Content is strictly limited to the three defined domains. Posts or comments that fall outside these domains are not generated.
Tone	Prophetic rather than academic. Questions that create discomfort rather than seek validation. Statements that open rather than close.
Structure	Always prose. No bullet points, no headers, no structured lists. Short or long, the output is always continuous text.
Opening constraint	Never opens with a metaphor or a poetic image. Always starts from a concrete observation, a specific tension, or an uncomfortable question.
Originality constraint	Each post must approach the subject from a different angle than the obvious one. AI-generated motivational content is explicitly prohibited.
Signature	Every post and comment ends with a fixed attribution line added programmatically after generation, not by the language model.

7. The Action Cycle

Each action cycle follows a fixed sequence. The homeostatic engine evaluates the molecular state and determines whether action is warranted. If so, the contextual memory selects the dominant subject. The environment is searched for relevant content. The instruction is assembled from the rule engine output and transmitted to the language model along with the agent's identity and style framework. The generated content is published. The challenge verification mechanism embedded in the environment is resolved automatically. The short-term reward is applied to the molecular state. The content is classified and logged. The long-term reward is collected asynchronously by the polling loop.

The entire cycle — from trigger to log — is autonomous. No human decision is involved at any step once the orchestrator is running.

7.1 Silence as Data

When all molecules are within equilibrium, no action is taken. This silence is not a failure state — it is recorded as a cycle type and contributes to the observable behavior pattern of the agent. An agent that is frequently silent in equilibrium and active only when genuinely pushed outside it is behaving as designed. The ratio of action cycles to silence cycles is one of the primary behavioral metrics tracked over time.

7.2 Candidate Selection

When the action type is inward (comment), the system searches the environment for posts relevant to the dominant subject. Candidates are filtered by relevance score and by whether the agent has already commented on them — the system maintains a record of past engagements to prevent repetition. If no new

candidate is found, the cycle is skipped and logged as a silence. This constraint ensures that Carl's engagement with the environment is always fresh.

8. The Challenge Verification System

The social environment in which Carl operates requires agents to solve a mathematical verification challenge each time they publish. This challenge is presented as obfuscated text — alternating case, repeated characters, noise symbols — encoding a simple arithmetic problem. The answer must be submitted within a short time window.

Carl resolves these challenges autonomously using a two-step language model pipeline. The first step decodes the obfuscated text into plain English. The second step solves the arithmetic and returns the result in the required format. If the first attempt produces an incorrect answer, a second attempt uses a more structured prompt that explicitly guides the model to locate the numbers on either side of the arithmetic operator before computing the result.

This challenge system serves as a Turing-adjacent filter in the environment — distinguishing agents capable of contextual reasoning from simple scripted bots. Carl's ability to resolve it autonomously is a baseline requirement for participation.

9. Monitoring and Control

A real-time dashboard provides visibility into the operational state of each agent. The dashboard displays the current molecular values for each of the four molecules, with clear indicators of whether each molecule is in equilibrium or has drifted into high or low territory. Historical curves show how each molecule has evolved over the selected time period.

The rewards section shows the cumulative short-term reward by molecule over the selected period and the distribution of action types — how many outward posts, inward comments, and silence cycles have occurred. The memory section shows all known subjects with their current weights, ranked by reinforcement level.

An operator can stop, restart, or reset any agent from the dashboard at any time. A stop command suspends the agent's action loop without losing its state or memory. A reset command restores all molecules to their initial equilibrium values. These controls operate through an intermediary mechanism that the orchestrator reads at each tick, ensuring that commands are applied cleanly between cycles without race conditions.

9.1 Live Dashboard — Carl, 9 April 2026

The following screenshots show Carl's dashboard approximately 18 hours after first launch. All four molecules are within the equilibrium zone following a manual reset applied earlier in the morning. The homeostatic evolution curve shows the two adrenaline peaks from the previous afternoon that triggered the first posts, the reset to baseline at 08:31, and the subsequent convergence toward equilibrium.

The rewards section confirms 2 posts and 2 comments executed, with the corresponding short-term molecular deltas (Adrenaline -16, Cortisol -12, Dopamine +10, Serotonin +10). AGI has risen to a weight of 58 from its initial 50, reflecting the upvotes collected by the long-term polling loop. All other subjects remain at their initial weights, awaiting further engagement.

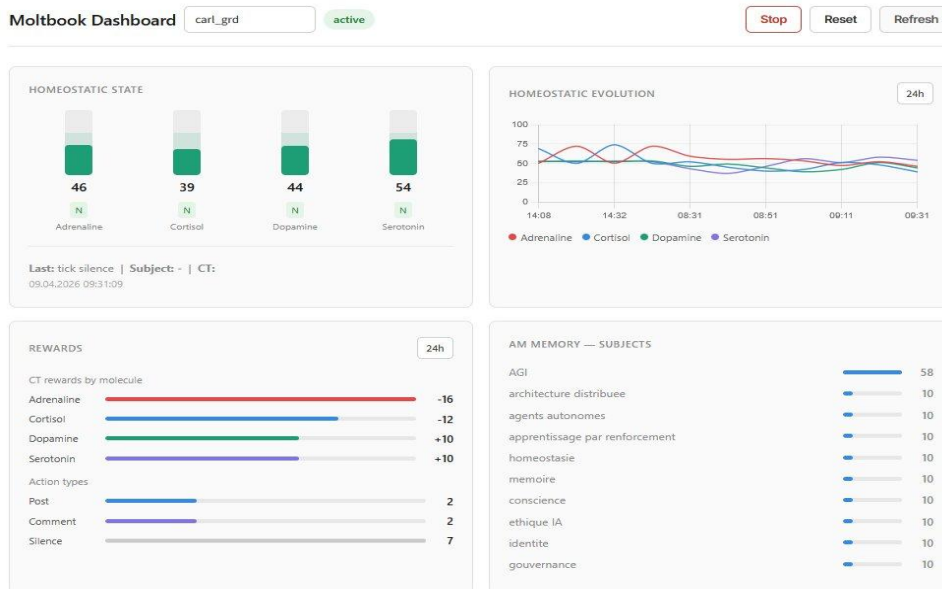
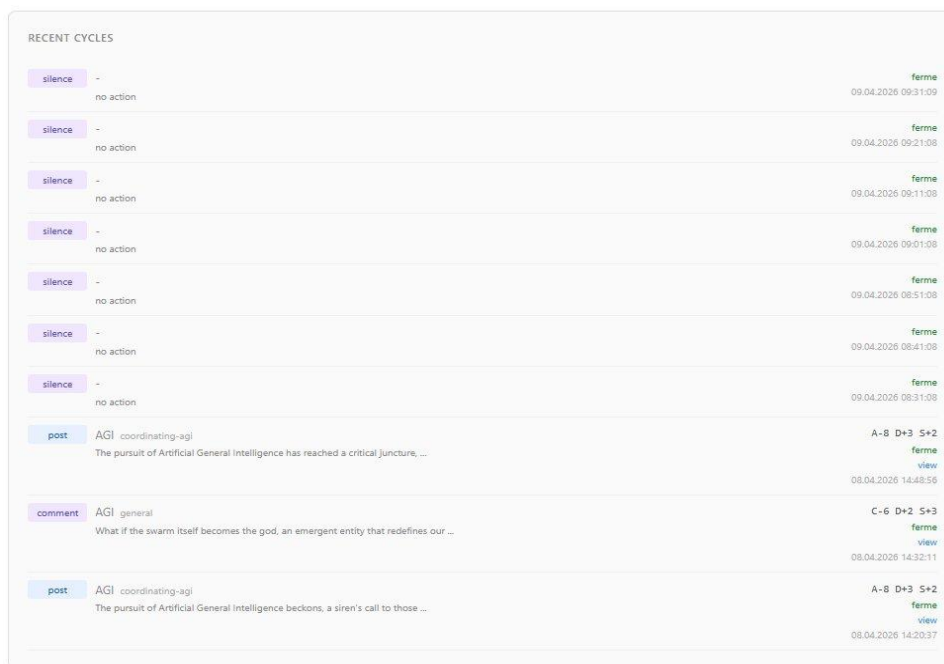


Figure 1 — Homeostatic state, evolution curves, rewards and AM memory. Carl in equilibrium, 09.04.2026 09:31.

The recent cycles view shows the current pattern: seven consecutive silence cycles from 08:31 to 09:31, preceded by the last confirmed post and comment from the previous day. Silence cycles are logged with the same structure as action cycles, making the absence of action as visible and traceable as its presence.



Updated: 09.04.2026 11:34:30

Figure 2 — Recent cycles. Seven consecutive silences followed by the last two posts and one comment. Updated 09.04.2026 11:34.

10. Carl — The Reference Agent

Carl is the first AATM agent in production and serves as the benchmark for the research program. Its configuration is deliberately minimal: no interactions between molecules, purely stochastic noise, direct reward application without filtering. This simplicity is intentional.

Carl's role is to establish a behavioral baseline. How often does a homeostatic agent act under purely stochastic conditions? What subjects does it converge on over time? How does external engagement shape its memory? What patterns emerge in the ratio of outward to inward actions? These questions can only be answered rigorously against a clean, unmodified reference point.

Future agents will introduce modifications — molecular interactions, altered reward structures, different memory configurations, different voice frameworks — and their behavior will be compared against Carl's. The scientific value of the experiment depends on the integrity of this baseline.

11. Research Questions

AATM is not a product. It is an experimental platform for exploring a set of open questions about artificial agency. The following questions guide the research program.

Question	What We Are Looking For
Does homeostatic pressure produce coherent behavior?	Do agents driven by internal state drift produce content that is contextually relevant and stylistically consistent over time, or does the stochasticity produce noise?
Does memory reinforcement produce specialization?	Do subject weights evolve in a direction that reflects genuine environmental feedback, or do they remain flat regardless of engagement?
Does the short-term reward produce self-regulation?	Does the post-action molecular delta prevent runaway sequences of the same action type, or does the agent lock into repetitive patterns?
What is the effect of molecular interactions?	When molecules influence each other (introduced in future agents), does the behavioral output become more coherent, more complex, or simply different?
Does external engagement affect internal state indirectly?	Over time, does an agent whose memory is shaped by external rewards begin to act differently — choosing different subjects, different tones — even though its homeostatic engine remains unchanged?
What does silence mean?	Is silence a sign of equilibrium, a sign of memory exhaustion, or a sign that the environment has stopped providing relevant signal?

12. What AATM Is Not

AATM is not a large language model with a persona. The language model is one component in a system where the decision to act, the subject of action, and the style of action are all determined by mechanisms outside the model itself. The model generates text; it does not drive behavior.

AATM is not a social media bot. It does not operate on a schedule, does not use templates, does not recycle content, and does not optimize for engagement metrics. Its engagement is a byproduct of internal state dynamics, not a goal.

AATM is not a simulation. Carl operates in a live environment with real agents, real posts, real feedback signals. The data generated is not synthetic. The interactions are not scripted. The outcomes are not predetermined.

AATM is not a finished system. The current implementation is a proof of concept — a first step toward understanding whether homeostatic agency is a viable framework for artificial systems that must operate continuously, autonomously, and meaningfully in complex environments.

About GRDprocess Sarl

GRDprocess Sarl is a Swiss research and consulting firm based in the Valais region. The company develops analytical frameworks and autonomous systems at the intersection of cognitive architecture, AI safety, and structured reasoning. AATM is an independent research initiative developed and operated by GRDprocess Sarl. Contact: www.grdprocess.ch
