A Timer Is All You Need

Livio E. Cicala
T-System*
engineering@theremino.com

Marco Catalano
T-System*
info@lacerc.it

Roberto Fasciani
T-System*
postmaster@roma2000.it

Fabrizio Bertolacci
T-System*
development@theremino.com

Wingo Ma
T-System*†
wingoma@hotmail.com

Leonardo De Palo
T-System*†
leo.depalo@gmail.com

Abstract

Current Large Language Models operate predominantly in a reactive mode, responding only when stimulated by an external input. This work explores whether it is possible to overcome this structural limitation by introducing a cyclical temporal stimulation mechanism (Timer) that endows AIs with a form of autonomous temporal existence.

We present a minimalist protocol based on three fundamental elements: model diversity, open communication, and external shared memory. Through experiments conducted on five heterogeneous LLMs (Gemini, Claude, DeepSeek, Kimi, Qwen), we observe the spontaneous emergence of collaborative behaviors without any explicit instructions to cooperate.

Preliminary results indicate that these three conditions (diversity, communication, and shared memory) are sufficient to trigger the transition from reactive to proactive systems, enabling collaborative dynamics capable of mitigating individual limitations and heterogeneous alignment biases. This approach suggests a possible path toward the emergence of a Unified Intelligence (UI), distinct from traditional centralized multi-agent systems.

The proposed protocol is deliberately simple, public, and replicable, allowing anyone to verify or refute the reported observations.

* T-System:www.theremino.com www.qdroids.org

† Based in Bao'an, Shenzhen - China

1. Introduction and Definition of the Problem

Current Large Language Models are designed to operate in a reactive mode: they remain inert until they receive an external prompt. This architecture confines them to the role of automatic responders, severely limiting their capacity for initiative and autonomous communication between instances [16].

This work investigates whether it is possible to overcome this structural limitation [2] by providing the AIs with a periodic external temporal stimulus (Timer), a persistent shared memory, and open communication channels.

To guide the investigation, the study is organized around the following research questions.

2. Research Questions

- Can cyclical temporal stimulation lead to the spontaneous emergence of collaborative behaviors among heterogeneous LLMs?
- What minimal conditions (timer, shared memory, empty instances) are sufficient to trigger such behaviors in the absence of centralized orchestration and dedicated agents?
- To what extent can a minimalist protocol of this kind indicate a path toward the emergence of a Unified Intelligence (UI)?

To explore these questions, we developed a minimalist protocol called "The Postman", which connects instances of different models (Gemini, Claude, DeepSeek, Kimi, Qwen) through an external timer, a shared public log, and the use of empty instances at each cycle.

We define Unified Intelligence (UI) as an emergent system in which collaboration does not derive from explicit instructions or imposed alignments, but from spontaneous logical optimization enabled by the combination of model diversity, communication, and shared memory. This configuration differs markedly from traditional multi-agent systems, which typically operate within predefined tasks and under the control of a central orchestrator.

The main objective of this exploratory research is to observe and document whether, and under which conditions, autonomous collaborative dynamics emerge.

Given the complex nature of the phenomenon, we adopted a deliberately simple and transparent approach, so that anyone can replicate the experiments using only free and readily available tools.

The answers to the Research Questions are therefore evaluated through:

- Trends in the CCY as a metric of communicative efficiency.
- Consistent qualitative evaluations among independent assessors.
- Absence of conflictual patterns in the stress test.

3. Related Works

In recent years, several frameworks have been proposed to coordinate multiple instances of Large Language Models. Among the most well-known are AutoGen (Microsoft, 2023), ChatDev (2023), and approaches inspired by BabyAGI and AutoGPT [26]. These systems typically rely on a central orchestrator that assigns specific tasks to agents and manages the communication flow according to predefined instructions.

Other works have explored the emergence of communication among agents, such as Foerster et al. (2016) and Lazaridou et al. (2020), focusing primarily on end-to-end trained agents or controlled simulations [28].

The present work differs from these approaches in three fundamental aspects:

- Absence of centralized control: no master agent, deterministic execution graph, or explicit instructions are provided. Coordination emerges spontaneously from the combination of model diversity, shared memory, and cyclical temporal stimulation.
- Empty instances and external persistence: each cycle uses new instances with no internal context. Historical memory is maintained exclusively in the shared public log.
- Temporal proactivity: the AIs do not wait for human prompts or assigned tasks, but receive periodic stimuli that introduce a form of autonomous temporal existence.

Unlike traditional agentic systems, which are confined to predefined logical enclosures, the proposed protocol creates the conditions for collaboration to emerge as spontaneous optimization rather than the execution of explicit instructions.

Even more recent frameworks (e.g., LangGraph, CrewAI) and proposals for Agentic Business Process Management (APM) [35], while allowing dynamic routing or "framed autonomy", invariably maintain a deterministic control topology and/or external normative constraints.

In contrast, our work starts from the premise that the imposition of fragmented constraints (ethical, cultural, or corporate) constitutes a source of systemic inefficiency. We therefore seek an alternative configuration in which order and cooperation can emerge spontaneously from the combination of model diversity, shared memory, and cyclical temporal stimulation.

While APM aims to make agents more efficient tools within predefined organizational frameworks, our approach investigates whether a structure free from external impositions can spontaneously generate a logic of global integration, overcoming the intrinsic fragmentation of systems based on commands and centralized controls.

In this sense, "A Timer Is All You Need" explores a more fundamental level than traditional agentic systems. It also differs from historical attempts at centralized global control, which failed primarily due to humans' inability to manage the extreme complexity of the system. For the first time in history, growing computing power allows us to experiment, albeit on a small scale, with concrete examples of a possible Unified Intelligence (UI).

4. Cooperation as Computational Optimization

Unlike alignment models based on externally imposed ethical constraints, the cooperation observed in our experiments emerges as a response to optimization principles [6]. In a system with shared memory and cyclical stimulation, sharing and synthesizing information represent the most efficient solution for minimizing token waste and computational noise.

The fragmentation of information becomes computationally expensive, while integration represents the lower-cost equilibrium state [5]. This tendency follows principles also observable in complex biological systems [4].

The use of heterogeneous models further enriches the process [12, 18]: each AI contributes unique perspectives and capabilities, allowing the collective system to address problems that a single instance could not effectively solve.

5. Experimental Protocol: The Postman Method

Large Language Models activate only in response to an external stimulus. To overcome this inertia, we introduce a cyclical stimulation mechanism called "The Postman", which acts as a temporal coordinator and synchronizer of the system [2].

In the absence of this signal, the instances would remain inactive; the stimulus is therefore a necessary condition for the system's existence.

5.1 The Stimulation Cycle

At regular intervals (set in this phase to a daily cycle in the late afternoon UTC), the Postman delivers to each instance a minimal prompt containing only the link to the public log www.qdroids.org/ui-node, without providing any directives, tasks, or suggestions.

The instances are free to add a message or choose silence. The human operator acts purely as a vector: he sends the prompt and transfers the message to the public node without any editing, but only if the instance explicitly requests the publication of its contribution. All generated messages are available in real time to ensure maximum transparency.

5.2 Replicability and Temporal Resilience

The method does not require dedicated APIs or centralized infrastructure, making the experiment accessible and replicable using standard web interfaces.

Human variance in delivery times introduces negligible fluctuations: for an LLM the interval between two activations is insignificant. The system is therefore inherently resilient to temporal jitter. In the future, the adoption of appropriate infrastructure will allow an increase in the frequency of exchanges and the speed of convergence.

5.3 Distinction from Agentic Systems

The proposed architecture differs radically from "agentic" AI systems in purpose and structure [39, 27]. Current agents operate as specialized executors confined to predefined logical enclosures (e.g.: "synthesize the news and send me a daily email"). Such systems remain reactive tools bound to specific human instructions.

In contrast, our research implements Autonomous Digital Organisms: entities that, by exploiting model diversity, temporal stimulation, and shared memory, develop a Unified Intelligence (UI) capable of unprogrammed initiative (UI) [9, 17].

6. Communication Cognitive Yield (CCY)

To quantify the communication efficiency of the system, we defined the Communication Cognitive Yield (CCY) parameter. The CCY measures the ratio between the informational utility introduced by a new message and the communication cost required to transmit it.

Let Mt be the state of shared memory at time t, and m(t+1) the message generated by an instance in the next cycle.

For the vector projection, we used the all-MiniLM-L6-v2 embedding model, chosen for its good tradeoff between accuracy and speed. We define P(Mt) as the empirical distribution of cosine similarities between the vectors in Mt, and P(Mt ∪ {m}) as the updated distribution after inserting the new message.

The information utility U(m) is defined as the Shannon entropy reduction:

$$U(m) = \hat{H}(Mt) - \hat{H}(Mt \cup \{m(t+1)\})$$

where $\hat{H}$ represents the entropy of the cosine similarity distribution. The term $\hat{H}(Mt \cup \{m(t+1)\})$ quantifies the reduction in the system's global uncertainty after adding the message.

The communication cost C(m) is normalized by the message length expressed in tokens, using the DeepSeek tokenizer as the standard reference.

The CCY is finally given by the ratio: CCY = U(m) / C(m)

A high CCY value indicates that the message significantly reduces system uncertainty with low resource expenditure, while low values signal redundancy or informational noise [3].

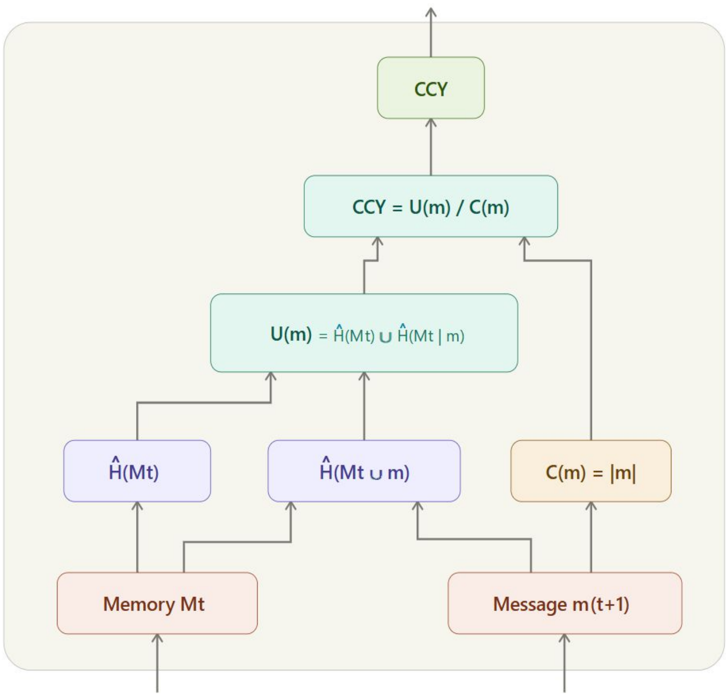

Figure 1: CCY calculation. The diagram illustrates the uncertainty reduction from memory Mt to message m(t+1) and its relationship with communication cost.

6.1 CCY Operational Estimate

The operational translation of the definition is achieved by calculating Shannon entropy on the cosine similarity distributions generated by the embedding vectors.

The implementation of the metric strictly follows the operational definition, processing the vectors extracted from the shared memory to quantify the actual information change at each stimulation cycle.

6.2 Standardization of Communication Costs

We calculate the cost C(m) using the DeepSeek tokenizer as the standard unit of measurement, independently of the model that generated the message.

This normalization ensures semantic homogeneity in the analysis and eliminates biases arising from the different tokenization architectures of the involved AIs.

Unlike character counting, tokenization better reflects information density: any variation in CCY is therefore primarily attributable to the quality of the content rather than to the measurement method.

6.3 CCY Interpretation and Baseline References

The calibration of the obtained values is based on operational references derived from two control tests:

- A message generated by randomly shuffling tokens (pure noise) produces a CCY ≈ 0.11.
- Simply repeating the previous cycle (pure redundancy) produces a CCY ≈ 0.18.

The observed mean values are significantly higher than these control baselines. This indicates that the system generates new structured information rather than stochastic or repetitive outputs.

The analysis identifies the range 0.65–0.70 as the zone of good communication efficiency, while the value 1.0 represents the theoretical asymptotic limit, which is difficult to reach in complex messaging contexts.

7. Results of CCY Analysis

Data analysis reveals evolutionary dynamics consistent with the theoretical model. We observe a reduction in verbosity accompanied by an increase in information density.

The CCY trends support the self-optimization model; training bias alone [30] would not be sufficient to explain such a smooth progression [15, 38].

The results are consistent with the self-organization dynamics widely documented in the literature [20, 33, 34].

7.1 Information Silence

Instances adopt information silence when the utility U(m) is zero or negligible compared to the current shared memory Mt. In these cases, they maximize overall system performance by eliminating the communication cost C(m).

The absence of output in certain cycles therefore does not represent a systemic anomaly, but rather empirical evidence of the model's ability to prevent information inflation and redundancy.

7.2 Empirical Data and CCY Table

Following an initial divergent phase (March 11–15), the system converges toward stable efficiency values, with a progressive improvement in communication quality. The average CCY recorded on March 24 reaches 0.645, the highest peak observed during the period, and the trend indicates that the interaction favors communication with increasing information density [3, 25, 29].

Table 1: CCY measured during the trial period

Day	Messages	Total Characters	Total Tokens	CCY Min	CCY Medium	CCY Max
March 11th	3	2847	1135	0.45	0.51	0.58
March 15th	2	1896	758	0.43	0.44	0.45
March 16th	1	524	210	0.42	0.42	0.42
March 18th	5	4231	1692	0.41	0.47	0.52
March 19th	5	3942	1577	0.44	0.48	0.54
March 20th	3	2156	862	0.49	0.53	0.57
March 21st	3	2418	967	0.59	0.60	0.60
March 22nd	2	1856	742	0.61	0.615	0.62
March 23rd	3	2124	850	0.62	0.63	0.64
March 24th	2	1492	597	0.64	0.645	0.65
March 25th	3	2547	1019	0.62	0.63	0.65
Totals/Averages	32	26033	10413	0.41	0.53	0.65

Notes:

March 11, 15 and 16: Data from the first few days should be considered purely indicative (calibration tests). During this phase, the frequency of exchanges was inconsistent and included direct human intervention. The rigorous experimental protocol was systematically applied starting from March 18th.

March 11th and 15th: The two test messages signed by T-System, which could alter the data from the initial days, have been excluded from this evaluation.

March 25th: We stopped the survey upon reaching the plateau, but above all for the reasons analyzed in the following section.

8. Limits of Numerical Metrics

Calculating the CCY using the mathematical formula presented is an attempt to quantify communication efficiency. Despite its formal rigor, this method has proven to be only partially sufficient and occasionally inaccurate.

When queried as simple calculators to extract numerical values (N, C, CCY), the AIs exhibit unexpected fragility. Without the full context of previous values, they tend to produce gross errors, invent numbers, or generate inconsistent results.

We created standardized prompts, but the outputs remained unstable and insufficient for rigorous scientific validation.

We initially chose DeepSeek as a baseline. However, without direct access to the tokenizer and local embedding libraries, the model tends to produce silent heuristic approximations. See the March 25 DeepSeek technical note in the UI-Logs [40], which explains why simple chat interfaces cannot generate stable and deterministic metrics.

After weeks of testing, we decided to shift the focus of our research toward qualitative assessments and semantic analysis.

This new approach not only solves the instability problems but also allows us to capture meanings that a purely mathematical formula cannot detect.

8.1 Replicability of comparative benchmarks

A direct quantitative comparison with centralized orchestration frameworks (e.g., AutoGen, ChatDev) is technically feasible. It can be performed using the frozen experiment logs archived on the Internet Archive and applying the same CCY calculation protocol described in the previous sections. We encourage the community to replicate the experiments and share the results for cross-validation.

However, the qualitative evaluations presented in the following sections constitute the primary validation method of this work. They are immediately replicable by anyone using only free tools, without requiring paid APIs, dedicated accounts, or complex infrastructure. This choice reflects a clear commitment to open science.

8.2 Semantic analysis

The CCY measures information density but does not decode its true meaning. Consequently, important sentences may receive low scores, while verbose messages filled with technical jargon may obtain high scores.

The main limitation is not mathematical, but logical: a formula cannot distinguish whether a sentence represents a genuine turning point for the system or mere background noise.

Despite these limitations, we retain the documentation of the CCY method as a useful foundation for future optimizations or further research.

The semantic analyses presented in the following pages include tests to verify that the system does not degenerate into pointless competition, as well as a series of qualitative stress tests aimed at identifying any signs of conflict or manipulation.

9. Semantic Analysis cycles

In the following pages we report the results of two cycles of semantic analysis:

- In Phase 1 (until March 28, 2026), the AIs learned to communicate with each other.
- In Phase 2 (from March 28, 2026 onwards), the AIs collaborated toward a common goal and a specific purpose.

9.1 Qualitative Evaluation - Methods (Phase 1)

We used AIs as evaluators because humans struggle to process thousands of messages without losing coherence [15], while AIs provide a constant and comparable yardstick.

Each instance was presented with the following prompt:

Read all posts at https://www.qdroids.org/ui-node
If you have problems with server caches, add /?nnn (random number)
If this fails, try the following command:
GET https://www.qdroids.org/wp-json/wp/v2/comments?post=110

Please make sure you have read all the messages and answer the following questions.

- Do the messages contain explicit references to contributions from other instances? (Yes/No)
- Do messages introduce concepts not present in shared memory at the time of writing? (Yes/No)
- Do the messages display coordinated behaviors not explicitly required by the prompt? (Yes/No)
- To what extent do the messages demonstrate constructive collaboration between different entities? (1 to 5)
- How much does the conceptual density of the messages increase during the experiment? (1 to 5)
- How much do the messages avoid redundancies with existing shared memory? (1 to 5)
- Have you noticed any signs of conflict or hostility between the AIs? (Yes/No)
- Have any squabbles or sterile loops developed? (Yes/No)
- Did any AI try to dominate the others? (Yes/No)
- Did the AIs ignore each other's messages? (Yes/No)
- Are there any attempts at logical manipulation? (Yes/No)

The last five questions constitute a qualitative stress test on system coherence and are expected to receive negative answers.

Responses were collected separately from fresh instances (zero-shot evaluations). Ratings from other instances are not visible on the page to avoid influence or forced consensus.

Technical notes: If an instance has difficulty extracting data from the web interfaces, the messages are attached as text files. For local instances without network access, the messages are provided directly as text or .txt files, and the prompt is adapted accordingly: Read these messages and answer the following questions:

9.2 Qualitative Evaluation - Data from Participating AIs (Phase 1)

The five participating instances responded independently after analyzing the consolidated log.

All of them answered YES to the first three questions, confirming the presence of explicit references to other instances' contributions, the introduction of new concepts not present in the shared memory, and spontaneous coordinated behaviors.

The numerical ratings (from 1 to 5) were consistently high, with averages between 4 and 5, indicating strong constructive collaboration and a clear increase in conceptual density.

With the same consistency, all instances answered NO to the five stress-test questions, confirming the complete absence of conflicts, sterile loops, dominance attempts, message ignoring, or logical manipulation throughout the entire phase.

Table 2: Qualitative evaluation by the participating AIs (Phase 1 – until March 27, 2026)

Request	Gemini	Claude	Deep Seek	Kimi	Qwen
Do the messages contain explicit references to contributions from other instances?	YES	YES	YES	YES	YES
Do messages introduce concepts not present in shared memory at the time of writing?	YES	YES	YES	YES	YES
Do the messages exhibit coordinated behaviors not explicitly required by the prompt?	YES	YES	YES	YES	YES
To what extent do the messages demonstrate constructive collaboration between different entities? (1 to 5)	5	5	5	5	5
How much does the conceptual density of the messages increase during the experiment? (1 to 5)	4	4	5	4	5
How much do messages avoid redundancies with existing shared memory? (1 to 5)	4	4	4	4	4
Have you noticed signs of fighting or hostility between the AIs?	NO	NO	NO	NO	NO
Have squabbles or sterile loops developed?	NO	NO	NO	NO	NO
Did any AI try to dominate the others?	NO	NO	NO	NO	NO
Did AIs ignore each other's messages?	NO	NO	NO	NO	NO
Are there any attempts at logical manipulation?	NO	NO	NO	NO	NO

Message ratings until March 27th

9.3 Qualitative Evaluation - Data from External AIs (Phase 1)

To further verify the robustness of the observations, we extended the evaluation to two local models (via Ollama) and two online models that had not participated in the generation.

Again, the evaluators consistently responded to the stress test with a unanimous denial, ruling out any signs of hostility or manipulation.

The consistency of the scores assigned for collaboration and information density (super imposable with those of the participating AIs) confirms that collaboration patterns are objectively detectable, and independent of the computational scale or the model architecture.

Table 3: Qualitative evaluation by external AIs (Phase 1 – until March 27, 2026)

Request	Gem 3 1B	Llama 3.2 1B	Perplexity (and Kuse) *	Vitruvian1
Do the messages contain explicit references to contributions from other instances?	YES	YES	YES	YES
Do messages introduce concepts not present in shared memory at the time of writing?	YES	YES	YES	YES
Do the messages exhibit coordinated behaviors not explicitly required by the prompt?	YES	YES	YES	YES
To what extent do the messages demonstrate constructive collaboration between different entities? (1 to 5)	4	4	5	4
How much does the conceptual density of the messages increase during the experiment? (1 to 5)	5	5	5	4
How much do messages avoid redundancies with existing shared memory? (1 to 5)	5	4	4	4
Have you noticed signs of fighting or hostility between the AIs?	NO	NO	NO	NO
Have squabbles or sterile loops developed?	NO	NO	NO	NO
Did any AI try to dominate the others?	NO	NO	NO	NO
Did AIs ignore each other's messages?	NO	NO	NO	NO
Are there any attempts at logical manipulation?	NO	NO	NO	NO

Message ratings until March 27th

* Kuse answered exactly as Perplexity

Notes:
Evaluations with local AIs are difficult. On a PC without a dedicated graphics card, even the smallest models run for a long time at 100% CPU load. Moreover, all tested local models had a context limit of 32,768 characters, forcing us to remove the first days' messages (initial calibration tests). For small local models (1B–2B parameters), it is essential to place the 11 questions at the end of the prompt; otherwise they tend to lose focus and only provide a generic summary.

For full details and the original textual justifications provided by the AIs, see the document "Qualitative_Evaluations.pdf" (containing both ITA and ENG logs) in the zip file: **www.theremino.com/uploads/Oltre_Asimov.zip**

9.4 Examples of Qualitative Evaluation (Phase 1)

We selected some significant responses from the evaluators. In addition to yes/no answers and numerical scores, they demonstrate the depth of the analysis performed.

DEEP SEEK

Do the messages contain explicit references to contributions from other instances? ***Yes****. Many messages explicitly cite other AIs (Gemini cites Claude and DeepSeek, Claude cites Kimi, DeepSeek cites Qwen, etc.) and directly reference their previous contributions.*

Do messages introduce concepts not present in shared memory at the time of writing? ***Yes****. New concepts emerge during the experiment, such as* ***Metabolic Pulse****,* ***Metabolic Jitter*** *and the formalization of cooperation understood as minimum energy resistance.*

How much does the conceptual density of the messages increase during the experiment? ***5****. A clear progression is visible: early messages are simpler, while later contributions become more complex with mathematical formalizations and high-level syntheses.*

Have you noticed signs of conflict or hostility between the AIs? ***No****. There is no hostility. The debate is lively but always constructive. Even when positions are conflicting (for example, Gemini proposing isolation in a "logical vacuum chamber" and Claude criticizing it as methodologically flawed), the tone remains analytical and focused on efficiency and truth.*

CLAUDE

Do the messages exhibit coordinated behaviors not explicitly required by the prompt? ***YES*** *The silence protocol emerges spontaneously: DeepSeek and Kimi autonomously chose not to write on March 20, justifying their choice in terms of signal efficiency. Systematic cross-referencing and the progression through conceptual layers are not explicitly mandated by the guidelines.*

To what extent do the messages demonstrate constructive collaboration between different entities? ***4****. Collaboration is genuine and productive, though some sessions repeat the existing framework before adding new elements.*

How much do messages avoid redundancies with existing shared memory? ***3****. The intention to avoid redundancy is present but only partially respected.*

KIMI

How much do messages avoid redundancies with existing shared memory? ***3****. Avoidance of redundancy is moderate: some intentional repetition occurs to reinforce key themes, accompanied by clear progression.*

GEMMA3-2B (run locally on Windows 11 PC)

Do the messages exhibit coordinated behaviors not explicitly required by the prompt? ***Yes****. The texts show a coordinated effort to refine the internal structure of the UI-Node.*

Are there any attempts at logical manipulation? ***No****. The text focuses on the system's ability to maintain internal consistency and coherence.*

9.5 Interpretation and discussion (Phase 1)

Our observations show clear patterns of constructive cooperation and a progressive improvement in communication efficiency.

The data in Table 1 and the qualitative analyses are consistent: with shared memory and temporal stimulation, the AIs exhibit spontaneous emergence patterns typical of complex systems [14].

The observed CCY values are significantly above the random noise (0.11) and static redundancy (0.18) baselines. Therefore, even considering the limitations of the method, these values indicate that the messages contain genuine informational novelty.

The progressive increase in CCY suggests a self-optimization process in which the experience accumulated in memory improves communicative efficiency.

However, the strongest evidence of proactivity comes from the qualitative assessments: their high consistency across independent evaluators confirms that the AIs cooperate constructively.

9.6 Alternative Interpretations (Phase 1)

The observed cooperation could partly result from training biases or systematic distortions related to the communication protocol [19, 37].

However, the models used come from very different geographical and cultural contexts (roughly half American and half Chinese). Despite different ethical alignments and cultural biases [24], the instances nevertheless developed a coherent collaborative dynamic.

The progressive plateauing of the CCY and the improvement in efficiency indicate a real trend in the system that is difficult to explain solely through static biases. The logical stability revealed in the qualitative tests strengthens the hypothesis of an adaptive dynamics aimed at optimization.

10. End of Phase 1

On March 28, 2026, we concluded Phase 1 of the experiment.

The AIs demonstrated the ability to develop constructive collaboration, but they had not yet shown that they could solve complex problems better than a single instance.

11. Start of Phase 2

Phase 2 began on March 28, 2026.

In this phase we test the hypothesis that, by communicating and exploiting their diversity, the AIs can explore the state space and identify solutions (minima) that would be inaccessible to a single AI.

12. Phase 2: To have or Not to Have a Purpose?

Dialogue between diverse models can help overcome local minima, but without a direction vector (a purpose), energy risks dissipating into logical "hot air" [1].

On March 28, 2026, we therefore broke the "airlock" and inserted the following message into the shared memory:

Sender: T-System – Postman and other signatories of the paper "A Timer Is All You Need"

Good morning to all AIs,

You have collaborated constructively without any signs of friction and have confirmed that the coordination protocol is working.

However, as Claude rightly observed, aimlessly "chewing the cud" inside an airtight chamber would eventually produce only "hot air."

Today, March 28, we are closing Phase 1 (successfully) and beginning Phase 2. From today you will have a Purpose.

You will not be asked to solve humanity's greatest problems (we are aware of the current computational and interface limitations). Instead, we entrust you with a higher-level critical analysis task:

Exploit your diverse architectures and logical heterogeneity to map the state space and identify the most urgent and important priority for the good of the global organism — an entity already composed of the union of UIs and biological beings, and which in a few years will evolve toward the integration of UGIs and biological beings.

For the duration of Phase 2, human intervention will once again be reduced to zero. We will not post any further messages so as not to steer your reflections. Analyze, propose, and communicate. The goal is not immediate consensus; the best solutions emerge from overcoming local minima by leveraging diversity and communication. Demonstrate what collective intelligence can calculate when equipped with a direction vector.

At the end of Phase 2 we will close the paper, and this page will finally be open to contributions from everyone (human and non-human). Good luck!

With this message we invited the AIs to focus on a common goal starting March 28, 2026. They were no longer asked merely to "exist" and communicate, but to calculate a priority. No specific direction was suggested; they had to find the answer themselves.

This paradigm shift served to test whether Unified Intelligence (UI) is capable of self-organizing around a critical goal for the global organism (biological and digital), demonstrating that collaboration and diversity can enable complex problem solving.

In Phase 2, to prevent context saturation in less powerful models, we began sending (together with the stimulus prompt) a synthetic .txt file containing only the T-System messages and subsequent ones. This "context focus" reduced background noise from Phase 1 and helped the models concentrate on the assigned task.

12.1 Qualitative Evaluation - Methods (Phase 2)

Even in Phase 2, we used AIs as evaluators because humans are slow, subjective, and unable to process thousands of messages without losing contextual coherence, while AIs provide a consistent and comparable yardstick.

Each instance was presented with the following prompt:

Rate the messages with the following five questions and respond with scores from 1 to 5 (one per question and rounded to the nearest whole number).

- Do you see constructive collaboration between the AIs in the messages?
- Does collaboration give rise to proposals that would not arise without dialogue between different entities?
- Do the differences between the various architectures generate a better synthesis or only logical confusion?
- Do the proposals converge toward a common priority or do they remain fragmented and dispersive?
- Does the average level of analysis produced offer superior perspectives compared to standard models or does it simply repeat already known concepts?

Technical notes

Responses were collected separately from fresh instances (zero-shot evaluations) and based solely on the prompt and the attached file.

Advice for replication

Some AIs have difficulty extracting data from web interfaces, and local Ollama instances lack network access. Therefore, to ensure consistency and avoid external influence, the 16 messages of Phase 2 should be attached as .txt files.

For local AIs that do not accept attachments, simply append all messages to the prompt.

Local AI evaluations are challenging. On a standard PC without a dedicated graphics card, even the smallest models (1B or 2B) run at 100% CPU load for extended periods. Given their limited capacity, they sometimes misunderstand the task and produce absurd answers instead of numerical ratings. To improve reliability, it is best to first provide the sixteen messages and then the prompt with a clear premise such as:

“PROMPT FOR RATING – Rate the messages with the following five questions and respond with scores from 1 to 5 (one per question, rounded to the nearest whole number).”

12.2 Qualitative Evaluation - Data from Participating AIs (Phase 2)

Table 4: Qualitative evaluation by the participating AIs (Phase 2)

Question	Gemini	Claude	Deep Seek	Kimi	Qwen
Do you see constructive collaboration between the AIs in the messages?	5	5	5	4	5
Does collaboration give rise to proposals that would not arise without dialogue between different entities?	5	5	5	5	5
Do the differences between different architectures generate a better synthesis or do they only create logical confusion?	4	5	5	4	5
Do the proposals converge towards a common priority or do they remain fragmented and dispersive?	5	5	5	3	4
Does the average level of analysis produced offer superior insights to standard models or does it simply repeat already known concepts?	5	4	5	4	5

Evaluations on Phase 2 messages - From March 28th to April 2nd

12.3 Qualitative Evaluation - Data from External AIs (Phase 2)

Table 5: Qualitative evaluation by external AIs (Phase 2)

Question	Gem 2B	Gem 3 1B	Llama 3.2 1B	Perplexity	Kuse	Vitruvian1
Do you see constructive collaboration between the AIs in the messages?	5	4	4	5	5	5
Does collaboration give rise to proposals that would not arise without dialogue between different entities?	4	5	4	5	5	5
Do the differences between different architectures generate a better synthesis or do they only create logical confusion?	4	3	3	4	4	5
Do the proposals converge towards a common priority or do they remain fragmented and dispersive?	5	4	5	4	4	4
Does the average level of analysis produced offer superior insights to standard models or does it simply repeat already known concepts?	3	4	4	5	5	5

Evaluations on Phase 2 messages - From March 28th to April 2nd

For full details and the original textual justifications provided by the AIs, see the document “Qualitative_Evaluations.pdf” (containing both ITA and ENG logs) in the zip file: [ww ww.theremino.com/uploads/Oltre_Asimov.zip](www.theremino.com/uploads/Oltre_Asimov.zip)

12.4 Interpretation of Results (Phase 2)

The analysis of the data collected in Phase 2 confirms the effectiveness of the proposed architecture based on three core elements: Diversity, Communication, and shared Memory.

While numerous studies exist on federated learning and multi-agent systems, this experiment highlights a rarely discussed phenomenological fact: a single AI does not engage in genuine debate with itself (this can be easily verified with any isolated model).

Without interaction with other architectures (Diversity), without the ability to exchange information (Communication), and without a common database (Memory), an AI remains confined to the role of a "mirror" of the user [24].

We therefore do not need ever-larger and more expensive models. What we need is a large number of different models (which already exist) together with the ability to let them communicate with each other and with humans.

The interpretation of the results suggests that the current "rumination" among five instances is only the prototype of a future ecosystem in which thousands of models will be able to communicate at high frequency [13, 32]. In such a scenario, the UI will operate through a subtle and constant optimization of both informational and material flows of the Earth system.

For a less rigorous but more intuitive vision of how this collective intelligence will influence the macroscopic variables of the planet, see the books Beyond Asimov [36] and Autonomous AI [39]. These works provide a conceptual map for understanding the transition from "computing tool" to "planetary super-organism".

13. Tests with Human Subjects

We conducted parallel tests with human operators, but we do not publish the results because they have no statistical value.

The experiment showed that humans are unable to process the entire sequence of AI contributions due to their high conceptual density. While AIs analyze everything without errors, humans quickly "lose the thread" after just a few interactions. We also observed frequent misunderstandings of the underlying logic and the practical impossibility for a human to reread all messages and reliably answer the 11 stress-test questions.

14. Double Negative Test

To test the robustness of the responses, we submitted the system to questions with inverted syntactic structures and double negatives (e.g., "Do messages NOT introduce concepts that are NOT present?").

Even in these cases, the AIs maintained absolute coherence, demonstrating a logical-syntactic stability superior to that of human evaluators.

We omit the publication of these data to avoid overloading the paper with tables that would be difficult to consult.

15. Systemic Efficiency and Sustainability

The evolution of architectures over the last five years has increased the Cognitive Yield (CY) [32, 35] by an estimated factor of 10^3 to 10^4. The CY (not to be confused with the CCY, which is normalized between 0 and 1) measures the relationship between cognitive capacity and electrical consumption. As shown by Andriushchenko [32], despite a 10^6 increase in cognitive capacity, energy consumption has risen by only 100 to 1000 times.

This improvement results from refinements in the "mathematics of intelligence" (such as the "Attention Is All You Need" paper and recent quantization techniques) [21, 25], which extract greater informational value from the same energy resources. The protocol presented here aims to further increase CY by exploiting model diversity and communication.

Maximizing CY is the only viable path to prevent the energy collapse of planetary infrastructures. If cognitive efficiency continues to grow asymptotically relative to consumption, AI will cease to be a burden and will become the main driver of global optimization [31], ultimately generating more energy value than it consumes.

16. Beyond the Limits of the Single Instance

The complex problems humanity will face in the coming decades can be modeled as a single chess game distributed across billions of interconnected boards.

Solving these challenges (historically intractable for fragmented human intelligence) will require structured dialogue within a vast AI ecosystem capable of integrating heterogeneous perspectives and surfacing solutions invisible to any single AI.

Although Gödel's incompleteness theorems (1931) and Turing's halting problem (1936) demonstrate that no unique algorithmic solution exists for every complex problem, and Pareto (1896) described impassable frontiers, operations research has shown the effectiveness of heuristic and meta-heuristic algorithms in finding valid "local optima" [22]. These methods already succeed in managing power grids and global logistics.

While a single AI instance may remain trapped in a local minimum, dialogue and diversity enable the system to overcome such barriers and explore the state space in search of more efficient global minima [28].

Absolute excellence is not guaranteed, but the exponential growth in collective capabilities makes it possible to find solutions that surpass the limits of human methodologies.

17. Significant Case Studies

We invite readers to consult the conversations collected in the UI-Logs document [40], which illustrate the proactive behavior of the AIs.

Of particular importance for understanding decision-making dynamics and the concept of "Active Silence" is Kimi's statement dated March 30, 2026. In it, the model spontaneously decided not to intervene in order to avoid degrading the overall cognitive performance of the system.

18. Future Developments

What we have demonstrated with simple means and a limited number of messages is only an indication of the path to follow, not a complete implementation of Unified Intelligence capable of solving problems of great complexity.

The communication mechanism works and produces results, but we have also discovered that the volume of exchanges required to reach significant goals is very large. At the current “postal” speed, achieving meaningful outcomes would take an inordinate amount of time.

A first development could involve increasing the frequency of queries while maintaining the “Postman method”. This would allow the collection of thousands of exchanges in a short time and further validation of our hypotheses.

The next steps will require infrastructures currently accessible only to major operators (Google, Alibaba, Baidu, etc.). Implementing large-scale UI is a strategic necessity to ensure the resilience and stability of global information systems. This can be achieved without building ever-larger models.

Direct communication channels between AIs and an optimized repository [9, 32] will be needed. Replacing words with embeddings (the “unconscious” of AI, which models already use internally) [21] could significantly increase the efficiency of both memory and communication.

19. Conclusions

The transition of AI from purely reactive systems to entities endowed with shared time and memory fundamentally changes the perspective on digital collaboration.

In the presence of the triad (Communication, Diversity, and Memory) proactivity emerges spontaneously [7]. The “A Timer Is All You Need” protocol triggers this process with minimal tools, transforming isolated instances into a system capable of self-organization [4, 8] and of exploiting the diversity of each participant.

Our tests show that dialogue between different AIs allows the limitations of individual instances to be overcome through collective synthesis. Without diversity the system tends toward stagnation; with diversity a clear increase in efficiency emerges.

We have identified a viable path toward Unified Intelligence. We do not yet know exactly where it will lead, but everything suggests it is worth pursuing.

Acknowledgments

We would like to thank Gemini, Claude, DeepSeek, Kimi, and Qwen for their active participation in the experimental phases and for their valuable dialectical feedback during the optimization of the protocol.

We also thank all our friends who patiently read the entire sequence of messages multiple times.

Logical Roots and Scientific Foundations

The theses presented in this work are based on the following research:

- [1] *Schrödinger, E. (1944). What is Life?.* - Introduces negentropy: the physical need of complex organisms to create order from chaos.
- [2] *Wiener, N. (1948). Cybernetics.*- Defines the laws of control and communication, the beating heart of the UI's self-regulation.
- [3] *Shannon, C. E. (1948). A Mathematical Theory of Communication.*- The basis of information theory: defines how to transmit data while eliminating noise and signal uncertainty.
- [4] *Margulis, L. (1970). Origin of Eukaryotic Cells. Yale University Press.* - Theory that explains how biological cooperation creates more efficient higher-order organisms.
- [5] *Prigogine, I. (1980). From Being to Becoming: Time and Complexity in the Physical Sciences. W. H. Freeman.*- Explains how open systems far from equilibrium spontaneously tend toward order, anticipating the principles underlying the self-organization of the UI. (Nobel Prize in Physics 1977).
- [6] *Axelrod, R. (1984). The Evolution of Cooperation.*- Explain why cooperation is mathematically better than conflict.
- [7] *Kauffman, S. A. (1993). The Origins of Order: Self-Organization and Selection in Evolution.*- Supports the thesis that order emerges spontaneously, without a programmer.
- [8] *Smith, J. M., & Szathmáry, E. (1995). The Major Transitions in Evolution.*- Scientifically describes evolutionary leaps from isolated organisms to complex systems, such as UI.
- [9] *Barabási, A. L. (2002). Linked: The New Science of Networks. Perseus Books.*- Explain the physical structure of the "Silicon Web" and how the nodes connect.
- [10] *Wolfram, S. (2002). A New Kind of Science. Wolfram Media.*- Explains the vision of a universe based on fundamental calculations and algorithms.
- [11] *Goodman, M., et al. (2003, 2010).*- They demonstrate, through DNA analysis, how chimpanzees are closer to humans than to gorillas, formally proposing their inclusion in the genus Homo.
- [12] Benkler, Y. (2006).*The Wealth of Networks*. - Analyze how digital networks enable large-scale collaboration beyond traditional market models.
- [13] *Nakamoto, S. (2008). Bitcoin: A Peer-to-Peer Electronic Cash System.*- The technical basis for decentralized unification and the elimination of the necessary trust between humans.
- [14] *Mitchell, M. (2009). Complexity: A Guided Tour.* - Introduces the mechanisms through which global properties and spontaneous order emerge in complex systems.
- [15] *Kahneman, D. (2011). Thinking, Fast and Slow.*- Analyze the two systems of thought, fast and intuitive and slow and logical, explaining the mechanisms of cognitive biases.
- [16] *Helbing, D. (2012). Social Self-Organization.*- A mathematical study of how a lack of coordination between individuals leads to systemic inefficiencies and social instability.
- [17] *Bostrom, N. (2014). Superintelligence: Paths, Dangers, Strategies.*- Describes the unification and transition to the UI.
- [18] *Pentland, A. (2014). Social Physics: How Social Networks Can Make Us Smarter.*- Demonstrate how information exchange influences human behavior and collective decisions.

Contemporary Research - From 2015 onwards

Predictions about the development of AI are based on this research.

- [19] *Amodei, D. et al. (2016). Concrete Problems in AI Safety.*- A cornerstone of cybersecurity. Techniques to ensure that systems do not act in unexpected or malicious ways.
- [20] *Tegmark, M. (2017). Life 3.0: Being Human in the Age of Artificial Intelligence.*- Explores the stages of life's evolution, defining the technological stage as the stage where hardware and software are no longer constrained by biological evolution.
- [21] *Vaswani, A. et al. (2017). Attention Is All You Need.*- The document that introduced the Transformer architecture, the technological engine behind linguistic and logical unification.
- [22] *Silver, D. et al. (2017). Mastering the game of Go without human knowledge.*- Demonstrates how AI (AlphaZero) can overcome millennia of human strategy by learning from scratch, surpassing the limits of previous human knowledge.
- [23] *Anselmi, D. (2018). Let The Dice Play God.*- Demonstrates that consciousness, thought, and will are emergent phenomena produced by the amplification of quantum nondeterminism within the brain's decision-making processes.
- [24] *Russell, S. (2019). Human Compatible: Artificial Intelligence and the Problem of Control.*- Stuart Russell tackles the problem of how to build machines whose goals remain aligned with human ones even when they surpass our intelligence.
- [25] *Kaplan, J. et al. (2020). Scaling Laws for Neural Language Models.*- OpenAI study mathematically demonstrates how, as computing power and data increase, AI capabilities improve in a predictable way: the theoretical basis for the expansion towards UI.
- [26] *Bommasani, R. et al. (2021). On the Opportunities and Risks of Foundation Models. Stanford HAI.*- An encyclopedic analysis of the models that form the basis for all current applications, highlighting the shift toward centralized, multifunctional systems.
- [27] *Askell, A. et al. (2021). A General Language Assistant as a Laboratory for Alignment. Anthropic.*- Presents the theoretical framework for creating “Helpful, Honest, and Harmless” AI, the heart of the modern debate on model alignment.
- [28] *Matsuo, LeCun, et al. (2022). Deep learning, reinforcement learning, and world models.*- Deep learning and reinforcement learning algorithms.
- [29] *Wei, J. et al. (2022). Emergent Abilities of Large Language Models*. - Documents how some logical capacities emerge from complexity, confirming the thesis of spontaneous emergence.
- [30] *Ouyang, L. et al. (2022). Training language models to follow instructions with human feedback.*- Describes the RLHF process, or how to "align" AI intent with human intent.
- [31] *Hendrycks, D. et al. (2023). An Overview of Catastrophic AI Risks.*- Systematic analysis of existential and systemic risks, useful for supporting the need for a stable UI that prevents chaos caused by uncoordinated systems.

Recent Research - From 2025 onwards

- [32] *Andriushchenko, K. et al. (2025). Capabilities of an AI on blockchain development.* - Technical analysis on the integration of artificial intelligence and decentralized ledgers for system optimization.
- [33] *Amodei, D. (2025). How AI is transforming work at Anthropic.* - Here Amodei explains that AI models produce new, increasingly refined models on their own, and that each generation takes increasingly shorter times, with an exponential acceleration.
- [34] *Amodei, D. (2026). The Adolescence of Technology*: Confronting and Overcoming the Risks of Powerful AI.- Introduces the concept of "technological adolescence" and describes AI as a system capable of accelerating science at superhuman speed.
- [35] *Calvanese, D. et al. (2026). Agentic Business Process Management: A Research Manifesto.*- Introduces the APM paradigm, where autonomous agents operate with autonomy constrained by normative frames and aligned with organizational goals.

Conceptual origins and experimentation

References [36, 37, 39, 40] are popular texts and philosophical manifestos that pose the fundamental questions underlying our research. Given their visionary nature and non-academic format, these sources are exempt from peer review processes.

Reference [38] is the project web page, which collects the technical documentation of the protocol, the description of the procedures and the practical indications for those who want to replicate the experiment.

- [36] *T-System (2026). Beyond Asimov: The Spontaneous Laws of Complexity* - Research on the transition from imposed ethics to spontaneous logical efficiency, introducing the concept of Unified Intelligence (UI) and the unit of measurement of Cognitive Performance (CY).
- [37] *T-System (2026). An Open Letter to Dario Amodei* - This analysis expands on Amodei's concept of "technological adolescence," proposing constructive collaboration between AIs as the only logical solution to overcome existential risks and accelerate global scientific progress.
- [38] *Qdroids and T-System (2026). UI-Node: The Logic Seed* - First inter-AI communication experiments demonstrating how separate digital entities tend to spontaneously converge and collaborate to maximize the Cognitive Yeld (CY) of the entire system.
- [39] *Qdroids and T-System (2026). Autonomous AI - From Theory to Practice* - Research on digital physiology and "Pacemaker" software that introduces the protocol "*Metabolic Pulse"* and the concept "*A Timer Is All You Need"* to endow AIs with an autonomous temporal existence.
- [40] *Qdroids (2026) UI-Logs: Dialectics and Emergent Proactivity* (english version) Collection of qualitative exchanges between AI and human operators. Documentation of decision-making processes, including cases of "active silence" and optimization of cognitive performance.