

Bica 2 Bica 2

BICA 2



Quick Submit



Quick Submit



DIRECTORES DE TESIS DOCTORAL

Detalles del documento

Identificador de la entrega

trn:oid:::1:3571216258

Fecha de entrega

15 may 2026, 10:11 p.m. GMT-5

Fecha de descarga

15 may 2026, 10:17 p.m. GMT-5

Nombre del archivo

MetaDNA_CSR_v14_final.pdf

Tamaño del archivo

4.7 MB

24 páginas

10.521 palabras

64.222 caracteres





1% Similitud general

El total combinado de todas las coincidencias, incluidas las fuentes superpuestas, para ca...




Filtrado desde el informe

► Bibliografía

Grupos de coincidencias


-  **7 Sin cita o referencia 1%**
Coincidencias sin una citación ni comillas en el texto
-  **0 Faltan citas 0%**
Coincidencias que siguen siendo muy similar al material fuente
-  **0 Falta referencia 0%**
Las coincidencias tienen comillas, pero no una citación correcta en el texto
-  **0 Con comillas y referencia 0%**
Coincidencias de citación en el texto, pero sin comillas

Fuentes principales

- 1%  Fuentes de Internet
- 0%  Publicaciones
- 0%  Trabajos entregados (trabajos del estudiante)

Marcas de integridad

N.º de alerta de integridad para revisión

-  **Caracteres reemplazados**
74 caracteres sospechosos en N.º de páginas
Las letras son intercambiadas por caracteres similares de otro alfabeto.

Los algoritmos de nuestro sistema analizan un documento en profundidad para buscar inconsistencias que permitirían distinguirlo de una entrega normal. Si advertimos algo extraño, lo marcamos como una alerta para que pueda revisarlo.

Una marca de alerta no es necesariamente un indicador de problemas. Sin embargo, recomendamos que preste atención y la revise.

Grupos de coincidencias

- 7 Sin cita o referencia 1%**
Coincidencias sin una citación ni comillas en el texto
- 0 Faltan citas 0%**
Coincidencias que siguen siendo muy similar al material fuente
- 0 Falta referencia 0%**
Las coincidencias tienen comillas, pero no una citación correcta en el texto
- 0 Con comillas y referencia 0%**
Coincidencias de citación en el texto, pero sin comillas

Fuentes principales

- 1% Fuentes de Internet
- 0% Publicaciones
- 0% Trabajos entregados (trabajos del estudiante)

Fuentes principales

Las fuentes con el mayor número de coincidencias dentro de la entrega. Las fuentes superpuestas no se mostrarán.

- 1** Internet
jmir.org <1%
- 2** Publicación
Md. Tofael Ahmed, Zannatul Mawa Koli, Azmain Mahtab Rahat, Taslima Akhter, U... <1%
- 3** Trabajos del estudiante
Sungshin Women's University <1%
- 4** Internet
johal.in <1%
- 5** Internet
www.mdpi.com <1%
- 6** Publicación
Sune Holm, Maria Serban. "Philosophical Perspectives on the Engineering Approa... <1%
- 7** Publicación
Zhao-tun An, Hai-lei Kou, Yan Sun, Yan-sheng Wang, Xi-xin Zhang, Jia-qing Lu, Gu... <1%

Meta-DNA: A Regulatory Architecture for Adaptive Metacognitive Control in Intelligent Agents

Manuel Caro¹, Yirlis Pineda¹, and Darsana Josyula²

¹*Universidad de Córdoba, Montería, Colombia*

²*University of Maryland, Department of Computer Science, Maryland, USA*

Abstract

We propose that effective metacognitive control requires three computational properties that existing cognitive architectures do not jointly provide: (1) modular, condition-sensitive activation of discrete regulatory strategies; (2) emergent compound behaviors arising from the co-expression of multiple regulators; and (3) experience-dependent modulation of regulatory sensitivity over time. We introduce Meta-DNA, a formal cognitive architecture that operationalizes these properties as a dynamic regulatory system. Metacognitive strategies are encoded as discrete regulatory units (meta-genes), each with an explicit activation potential function and threshold. State-monitoring units (transcription factors) detect epistemic and cognitive signals and contribute to gene activation. When multiple meta-genes are co-expressed, compound metacognitive pathways are triggered, producing adaptive behaviors that no single regulator could generate independently. An epigenetic mechanism modulates regulatory thresholds based on accumulated experiential history. The architecture is specified as a discrete-time dynamical system and implemented in a controlled simulation. Across 30 independent runs per condition, Meta-DNA agents ($\theta = 0.50$) achieve a mean task accuracy of 0.731 (95% CI [0.726, 0.736]) versus 0.575 for non-metacognitive baselines (Cohen's $d = +8.25$, $p < 0.001$). Ablation studies confirm that removing `g_UncertaintyMonitor` produces the largest performance drop (d falls from +8.25 to +3.06), establishing its centrality to epistemic stress regulation. The biological analogy to gene regulatory networks (GRNs) provides the formal template for this architecture—not as a metaphor, but as a proven computational structure for achieving modular, composable, context-sensitive regulation of complex system behavior.

Keywords: metacognition, cognitive architecture, self-regulation, adaptive control, introspective reasoning, regulatory dynamics, epistemic adaptation, metareasoning

1. Introduction

Metacognition—the capacity of a cognitive system to monitor, evaluate, and regulate its own processing—is foundational to adaptive intelligent behavior (Flavell, 1979; Nelson & Narens, 1994). Formally, metacognition subsumes at least three distinct computational demands: the ability to detect when current reasoning strategies are failing or suboptimal (monitoring); the ability to select and deploy corrective strategies in response to internal state signals (regulation); and the ability to do both efficiently, so that the cost of metacognitive intervention does not exceed its benefit (Russell & Wefald, 1991; Cox, 2005). Satisfying all three demands simultaneously requires a control architecture that is simultaneously modular—so that individual metacognitive functions can be independently activated—composable—so that combinations of functions produce compound adaptive behaviors—and temporally adaptive—so that the sensitivity of metacognitive responses evolves with experience. No current cognitive architecture provides all three properties within a single formally specified framework.

Existing cognitive architectures address parts of this problem but not all of it. SOAR (Laird, 2012) handles monitoring through impasse-driven subgoalings but does not support compositional

regulatory logic: when multiple metacognitive conditions co-occur, SOAR has no mechanism to produce behaviors that emerge specifically from their combination. ACT-R (Anderson et al., 2004) enables procedural metacognition through conflict resolution and memory monitoring, but its meta-level control parameters are largely static and not updated by experience. CLARION (Sun, 2007) introduces explicit meta-cognitive subsystems but without formal co-activation conditions between regulatory units. MIDCA (Cox & Raja, 2016) provides structured feedback cycles between meta-level diagnosis and object-level reasoning, yet its regulatory logic is procedural rather than compositional. More fundamentally, none of these architectures model the temporal adaptation of metacognitive sensitivity: the property that an agent's responsiveness to a given type of cognitive event should itself change as a function of whether metacognitive intervention in response to that event type has historically been beneficial or costly. This gap leaves metacognitive architectures unable to develop the kind of domain-specific regulatory profiles that are a hallmark of expert cognition (Efklides, 2011).

This paper addresses the gap by introducing Meta-DNA, a cognitive architecture whose central claim is that the three required properties—modularity, composability, and temporal sensitivity adaptation—can be simultaneously achieved by formalizing metacognitive control as a regulatory activation system. The key insight is that gene regulatory networks (GRNs) in molecular biology already solve an analogous problem: how to produce modular, composable, context-sensitive control of complex system behavior from a compact encoding of discrete regulatory units (Alon, 2007; Jacob & Monod, 1961). Meta-DNA does not use GRNs as a metaphor. It uses the formal structure of GRNs—condition-sensitive activation potentials, threshold-gated expression, co-expression-triggered compound pathways, and experience-dependent epigenetic modulation—as a precise computational template for specifying metacognitive control. Each structural element of the biological system maps to a specific, formally defined component of the cognitive architecture, as detailed in Section 3. The biological analogy is therefore the mechanism of implementation; the theoretical contribution is the cognitive architecture and the formal properties it provides.

The central theoretical claim is precise: metacognitive control is computationally equivalent to a regulatory activation problem, and the formal structure of GRNs is a correct—not merely analogous—solution to that problem. This claim has testable consequences. An architecture built on this structure should exhibit (a) condition-specific activation of individual metacognitive functions, (b) emergent compound behaviors that arise only when multiple functions are co-expressed, and (c) progressive adaptation of regulatory sensitivity across episodes. We test all three predictions empirically through ablation studies that isolate each structural component of the architecture.

We specify Meta-DNA as a discrete-time dynamical system, implement it in a controlled simulation environment with three task regimes of increasing adversarial difficulty, and evaluate it across 30 independent runs per condition. The optimized architecture achieves mean task accuracy of 0.731 (95% CI [0.726, 0.736]) versus 0.575 for non-metacognitive baselines (Cohen's $d = +8.25$, $p < 0.001$). Ablation studies show that removing the `g_UncertaintyMonitor` gene reduces the effect from $d = +8.25$ to $d = +3.06$ —the largest single-component loss in the study—confirming that the epistemic monitoring function is the primary driver of performance gains. The regulatory components collectively account for the full effect: no single component is redundant.

The main contributions of this work are:

- A theory of metacognitive control as a regulatory activation problem, establishing that the computational properties required for principled adaptive metacognition—modularity, composability of co-expressed regulators, and temporal adaptation of sensitivity—are jointly provided by a GRN-structured formal architecture.
- A discrete-time formal model specifying five meta-genes, four transcription factors, and five co-expression-triggered metacognitive pathways, with explicit activation potential functions ϕ_i ,

sensitivity functions ψ_j , decay factors δ_j , and pathway transformation operators η_k —making every regulatory parameter independently inspectable and testable.

- An epigenetic modulation mechanism $\chi_i(t) = \alpha f(st) + (1-\alpha)\chi_i(t-1)$ that implements experience-dependent adaptation of regulatory sensitivity within the formal architecture, operationalizing the development of domain-specific metacognitive profiles without external retraining.
- An empirical evaluation across 30 independent runs per condition showing that Meta-DNA ($\theta = 0.50$) achieves task accuracy of 0.731 vs. 0.575 for baselines ($d = +8.25$, $p < 0.001$), and that removing `g_UncertaintyMonitor` produces the largest performance loss in the ablation (d drops from $+8.25$ to $+3.06$), confirming the distinct and non-redundant functional role of each architectural component.

2. Related Work

2.1 Metacognition in Cognitive Systems Research

Metacognition has been studied extensively as a component of intelligent behavior encompassing self-monitoring, self-evaluation, and self-regulation (Flavell, 1979; Nelson, 1992). In the domain of artificial intelligence, Cox (2005) provides a foundational overview of metacognition in computation, identifying introspective monitoring, goal-driven regulation, and strategy selection as core metacognitive functions. Russell and Wefald (1991) formalize the metareasoning problem as the optimal allocation of deliberation effort, demonstrating that rational agents must reason about the value of computation itself.

Recent metareasoning frameworks have extended this work to characterize stopping criteria, resource-bounded reasoning, and epistemic assessment under uncertainty (Svegliato et al., 2020). These formalizations motivate the design of Meta-DNA's meta-gene functions and pathway activation logic, which operationalize similar resource-allocation and introspective decisions in a modular regulatory framework.

2.2 Cognitive Architectures and Meta-Level Control

Traditional cognitive architectures implement metacognitive functionality through explicit rule-based layers. SOAR (Laird, 2012) uses impasse-driven subgoaling to handle failures in object-level reasoning, triggering meta-level problem spaces when goal achievement fails. ACT-R (Anderson et al., 2004) incorporates metacognitive processes through declarative memory monitoring and conflict resolution among production rules. CLARION (Sun, 2007) extends this with explicit meta-cognitive subsystems governing learning and strategy selection layered over implicit subsymbolic processing.

MIDCA (Cox & Raja, 2016) represents a more recent advancement, formalizing a two-tiered architecture in which meta-level processes diagnose and guide object-level reasoning through structured feedback cycles. The MAPE-K loop (Kephart & Chess, 2003) provides a complementary engineering framework for autonomous systems, specifying Monitor-Analyze-Plan-Execute cycles that govern adaptive behavior. While these architectures capture important aspects of meta-level control, they rely primarily on procedural mechanisms and lack compositional, structured representations of regulatory activation logic.

2.3 Biologically Inspired Regulatory Models

Gene regulatory networks (GRNs) have been employed in developmental robotics and artificial life to model adaptive control in complex, dynamic environments (Kitano, 2002). These systems demonstrate that modular, condition-sensitive regulatory mechanisms can produce flexible adaptive

behaviors from compact encodings. Alon (2007) provides formal analyses of regulatory circuit motifs in biological systems—including feedforward loops, oscillators, and co-expression networks—that inspire the pathway-level architecture of Meta-DNA.

Epigenetic regulation, in which gene expression is dynamically modulated by experience-dependent signals without changes to the underlying sequence (Zhong et al., 2008), provides the conceptual basis for Meta-DNA's adaptive threshold modulation. This mechanism permits long-term modulation of metacognitive sensitivity based on accumulated performance history, an extension not addressed by prior cognitive architectures.

2.4 Neural-Symbolic and Hybrid Approaches

Neural-symbolic systems have sought to combine explicit symbolic reasoning with subsymbolic learning to support flexible cognitive control (Garcez et al., 2019). While effective in constrained domains, these architectures often sacrifice interpretability and modularity. Meta-DNA offers a complementary approach that achieves interpretability through a transparent regulatory specification while retaining biological plausibility and compositional structure.

In summary, while prior architectures provide useful foundations for meta-level control, no existing framework combines structured regulatory dynamics, formal compositional pathway activation, and epigenetic modulation within a unified formal specification. Meta-DNA fills this gap by introducing these properties in a computationally tractable and cognitively interpretable architecture.

3. Meta-DNA Architecture: Formal Specification and Regulatory Dynamics

3.1 Biological Analogy and Computational Mapping

The design of Meta-DNA is motivated by three cognitive requirements established in Section 1: modularity of metacognitive functions, composability through co-expression, and temporal adaptation of regulatory sensitivity. We argue that gene regulatory networks (GRNs) in molecular biology constitute a natural formal solution to these requirements, because GRNs already solve an equivalent problem in a different domain: how to produce modular, composable, context-sensitive regulation of complex system behavior from a compact encoding of discrete regulatory units (Jacob & Monod, 1961; Alon, 2007). The formal correspondence is not analogical but structural. In a GRN, transcription factors detect biochemical signals and gate gene expression; expressed genes produce functional outputs; co-expression of multiple genes triggers compound pathway effects; and epigenetic marks modulate expression thresholds over time. In Meta-DNA, transcription factors detect epistemic and cognitive signals and gate meta-gene expression; expressed meta-genes modulate cognitive core functions; co-expression triggers compound metacognitive pathways; and experiential history modulates activation thresholds. Figure 1 illustrates this structural mapping. Table 1 formalizes each mapping in tabular form. This structural correspondence means that we adopt the same formal architecture as GRNs, which motivates the expectation of analogous properties—modularity, composability, and regulatory stability—in the cognitive domain. Whether these properties hold is established empirically in §5 and formally characterised in §3.7; they are not assumed from the biological parallel.

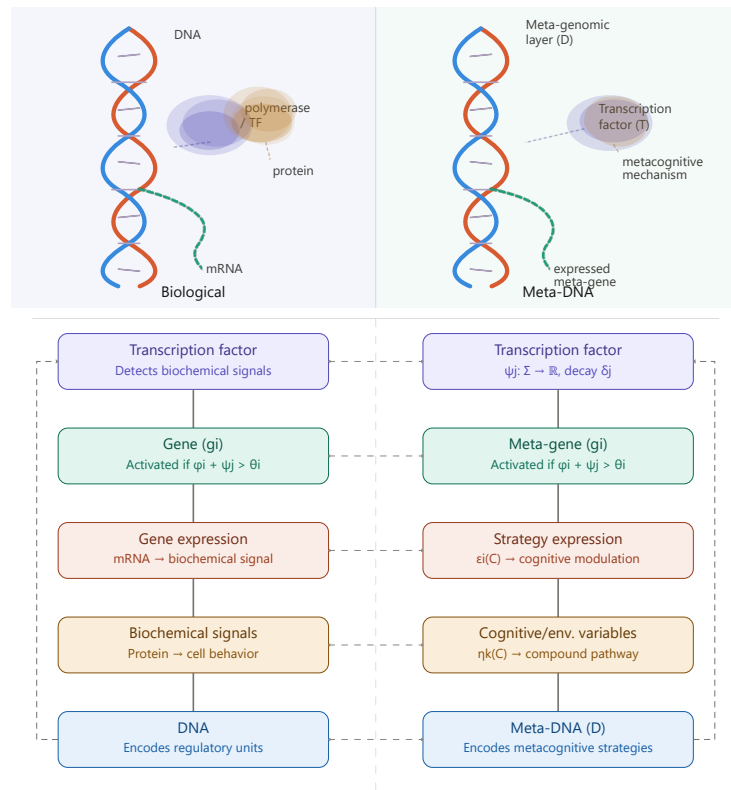


Fig. 1. Structural correspondence between gene regulatory networks (GRNs) and the Meta-DNA cognitive architecture. Each row represents one level of the regulatory hierarchy; horizontal dashed arrows (\equiv) indicate formal structural equivalence — not analogy — between biological and cognitive components. The epigenetic trace $\chi_i(t)$ (bottom row, pink) extends the GRN template with an experience-dependent threshold modulation mechanism absent from standard GRN formulations. Vertical arrows denote regulatory information flow within each column; curved dashed arcs indicate the closed regulatory feedback loop. Property labels [P1]–[P4] refer to §3.7.

Table 1. Mapping of biological transcription components to Meta-DNA architecture

Biological Concept	Meta-DNA Equivalent	Functional Role
DNA / Genome	Meta-Genomic Layer (D)	Encodes metacognitive strategies as latent regulatory units
Transcription Factor	Transcription Factor (T)	Detects epistemic/cognitive signals; regulates meta-gene expression
mRNA	Expressed Meta-Gene	Intermediate regulatory signal triggering cognitive modulation
Protein	Metacognitive Pathway Effect (η_k)	Functional transformation applied to the cognitive core
Epigenetic Mark	History Trace $\chi_i(t)$	Experience-dependent modulation of regulatory thresholds

3.2 Agent State and Self-Model

The agent's complete internal state at time t is a vector $st = (ut, lt, pt, et, zt) \in \Sigma \subseteq [0,1]^5$, where ut is epistemic uncertainty, lt is cognitive load, pt is task performance, et is energy level, and zt is satisfaction. This vector constitutes the observable surface of an explicit Model of the Self (MoS)—a structured self-representation that the metacognitive regulatory layer reads to decide which meta-genes

to express. The MoS is not an auxiliary data structure; it is the sole input to all transcription factor sensitivity functions ψ_j and all meta-gene activation potential functions ϕ_i . No regulatory decision in Meta-DNA is made without passing through the MoS.

3.2.1 Formal Structure of the MoS

The MoS integrates four categories of metacognitive content, following the knowledge taxonomy of Cox (2005):

- Declarative component D: encodes the current task regime (static, dynamic, or adversarial), observed difficulty and noise levels, and the binary outcome of the most recent decision. D provides the context that calibrates what constitutes “high” or “low” values of the continuous state variables.
- Procedural component P: encodes the set of currently available cognitive operations—in the simulation, the six modules of the cognitive core C (perception, memory, planning, reasoning, learning, action). This component informs `g_StrategyRevision` about which alternative strategies are available before triggering a revision.
- Conditional component C: encodes context-strategy associations accumulated over prior episodes—specifically, which combinations of (ut, lt) have historically triggered productive metacognitive responses versus incurred overhead without benefit. This component is operationalized by the epigenetic traces $\chi_i(t)$, which function as a continuously updated conditional knowledge store indexed by gene identity.
- Historical component H: the five continuous state variables (ut, lt, pt, et, zt) themselves constitute the running historical trace. Each variable is updated at every timestep via an exponential moving average toward the current environmental signal, making the MoS a compressed, recency-weighted record of the agent’s cognitive history rather than a snapshot of a single moment.

3.2.2 State Update Rules

Each of the five MoS components is updated exactly once per regulatory cycle (Algorithm 1, Step 7) by a deterministic rule parameterized by the environmental triple (dt, nt, bt) and the object-level decision outcome (correctness ct, energy cost kt). Equations (1)–(5) give the update rules for each component; the `clip(·)` operator enforces the domain constraint $st \in [0,1]^5$:

$$ut+1 = \text{clip}(ut + \Delta u + \xi_u, 0, 1) \quad (1) \quad \Delta u = (-0.03 \text{ if } ct \text{ else } +0.05) + 0.10 \cdot nt + 0.15 \cdot bt$$

$$lt+1 = \text{clip}(lt + \Delta l + \xi_l, 0, 1) \quad (2) \quad \Delta l = 0.05 \cdot dt - 0.02 + 0.20 \cdot bt$$

$$pt+1 = \text{clip}(0.92 \cdot pt + 0.08 \cdot p^* + \xi_p, 0, 1) \quad (3) \quad p^* = 0.85 \text{ if } ct \text{ else } 0.30$$

$$et+1 = \text{clip}(et - kt + 0.008 + \xi_e, 0.05, 1) \quad (4) \quad kt = dt \cdot 0.022 + U(0, 0.012)$$

$$zt+1 = \text{clip}(0.90 \cdot zt + 0.10 \cdot (0.6 \cdot pt + 0.4 \cdot (1-ut)) + \xi_z, 0, 1) \quad (5)$$

where $\xi_k \sim \mathcal{N}(0, \sigma_k)$ are i.i.d. Gaussian process noise terms with $\sigma_k \in [0.005, 0.015]$, independent across state dimensions. Note: ξ denotes process noise throughout §3.2; the symbol ϵ_i (used elsewhere in the paper) denotes the cognitive modulation function of meta-gene g_i ($\epsilon_i: C \rightarrow C$) and is a distinct quantity. The five update rules implement qualitatively different dynamics. Uncertainty (Eq. 1) is outcome-sensitive and perturbation-reactive: it decreases on correct decisions and spikes on perturbation events, making it the primary epistemic stress signal for `TF_UncertaintyMonitor`. Cognitive load (Eq. 2) accumulates with task difficulty and perturbation, with a slow baseline decay (−0.02 per step) modeling natural attention recovery between peaks. Performance (Eq. 3) is an exponential moving average ($\alpha=0.08$) toward a decision-outcome target, implementing a slow confidence signal that resists individual-trial noise; the asymmetric targets ($p^*=0.85$ correct, 0.30 incorrect) produce faster recovery than degradation. Energy (Eq. 4) depletes by task cost kt per step

with a fixed background recovery of 0.008, modeling a bounded cognitive resource with a floor at 0.05 preventing full depletion. Satisfaction (Eq. 5) integrates performance and epistemic clarity as a composite signal; its weighted combination ($0.6 \cdot pt + 0.4 \cdot (1 - ut)$) encodes the intuition that a highly uncertain but high-performing agent is less satisfied than a certain, lower-performing one — providing TF_StrategyRevision with a scalar signal for detecting prolonged suboptimality without requiring explicit per-trial error logging.

After the environmental update (Eqs. 1–5), each expressed meta-gene g_i applies its cognitive modulation function $\epsilon_i: C \rightarrow C$ to the cognitive core, which may alter state components of s^t within the same cycle (e.g., $g_UncertaintyMonitor$ reduces ut by 0.008 per step when expressed). The post-modulation state s^t is the final MoS value used as input at $t+1$. This creates a closed within-cycle feedback loop: the MoS state drives regulatory activation in Algorithm 1 steps 1–6, and the resulting modulations update that same state at step 7, making every regulatory decision self-consistent.

3.2.3 Concrete MoS States from the Simulation

Table 3 illustrates three representative MoS states drawn from the simulation (seed 7, Meta-DNA $\theta = 0.50$). Each row shows a qualitatively distinct regulatory situation and the downstream regulatory response it produces. These examples demonstrate that the MoS functions as a genuine cognitive self-model: different states correspond to recognizably different metacognitive situations, trigger different regulatory responses, and reflect different stages of the agent’s experiential history through the χ_i traces.

Table 3. Concrete Model of the Self states at three simulation steps (seed 7, $\theta = 0.50$)

Attribute	Step 50 — Static, settled	Step 334 — Regime transition	Step 500 — Dynamic, stable
State s^t	$u=0.248$ $l=0.443$ $p=0.855$ $e=0.074$ $z=0.832$	$u=0.588$ $l=0.750$ $p=0.737$ $e=0.039$ $z=0.832$	$u=0.400$ $l=0.949$ $p=1.000$ $e=0.036$ $z=0.911$
TF signals	UncMon=0.148 TimeA=0.238 EffB=1.000 StratR=0.026	UncMon=0.565 TimeA=0.608 EffB=1.000 StratR=0.000	UncMon=0.391 TimeA=0.715 EffB=1.000 StratR=0.000
Expressed genes	TimeAlloc, EffortBudget, PerfMonitor	TimeAlloc, EffortBudget, PerfMonitor	UncMon, TimeAlloc, EffortBudget, PerfMonitor
Activated pathways	PathwayTimeManagement	PathwayTimeManagement	TimeManagement, CognitivePlanning, IntrospectiveReporting
χ_i traces (cond. component)	UncMon=0.408 StratR=0.352 TimeA=0.747 EffB=1.000 PerfM=0.609	UncMon=0.168 StratR=0.201 TimeA=0.606 EffB=1.000 PerfM=0.604	UncMon=0.705 StratR=0.423 TimeA=1.000 EffB=1.000 PerfM=0.551
Metacognitive reading	Low uncertainty, energy-critical: energy conservation pathway active, no epistemic intervention needed	High uncertainty + load spike from perturbation: TF_UncMon active (0.565) but χ_{UncMon} still low (0.168) — uncertainty not yet recognized as chronic	High load, moderate uncertainty: χ_{UncMon} now elevated (0.705) — 166 steps of dynamic exposure have sensitized epistemic monitoring; three pathways co-active

The contrast between steps 334 and 500 is particularly informative. At step 334, the first perturbation of the dynamic regime has fired: uncertainty spikes to 0.588 and cognitive load to 0.750. TF_UncertaintyMonitor activates strongly (0.565), but the epigenetic trace χ_{UncMon} is still low (0.168)

because the agent has had only 334 steps of predominantly low-uncertainty experience. The MoS “reads” this as an acute stress episode rather than a chronic pattern, and the regulatory response is conservative—only PathwayTimeManagement fires. By step 500, the same uncertainty level (0.400, lower than at step 334) produces a qualitatively different regulatory outcome: χ_{UncMon} has risen to 0.705 through 166 steps of dynamic-regime exposure, making $g_{\text{UncertaintyMonitor}}$ now readily expressible. Three pathways fire simultaneously. This is the epigenetic sensitization process operationalized: the conditional component of the MoS has updated to recognize moderate uncertainty in a high-difficulty regime as a signal warranting compound epistemic intervention, not merely time management. The MoS has learned, within its own regulatory history, what that combination of signals means.

3.3 Architectural Components

Meta-DNA comprises four integrated components that collectively instantiate a structured metacognitive control loop (Figure 2, full schematic):

- Cognitive Core C: the object-level processing layer comprising modules for perception, memory, planning, reasoning, learning, and action selection.
- Meta-Genomic Layer D: a structured set of meta-genes $\{g_i\}$ encoding latent metacognitive regulation strategies. Each meta-gene $g_i = (\phi_i, \epsilon_i, \theta_i)$ consists of an activation potential function $\phi_i: \Sigma \rightarrow \mathbb{R}$, a cognitive modulation function $\epsilon_i: C \rightarrow C$, and an expression threshold $\theta_i \in \mathbb{R}$.
- Transcription Factors T: a set of state-monitoring units $\{t_j\}$. Each $t_j = (\psi_j, B_j, \delta_j)$ includes a sensitivity function $\psi_j: \Sigma \rightarrow \mathbb{R}$, a regulated gene subset $B_j \subseteq D$, and a decay factor $\delta_j \in \mathbb{R}^+$.
- Metacognitive Pathways M: compound regulatory operators $\{m_k\}$. Each $m_k = (R_k, \eta_k)$ activates when all meta-genes in the required set $R_k \subseteq D$ are co-expressed, applying a composed transformation $\eta_k: C \rightarrow C$ to the cognitive core.

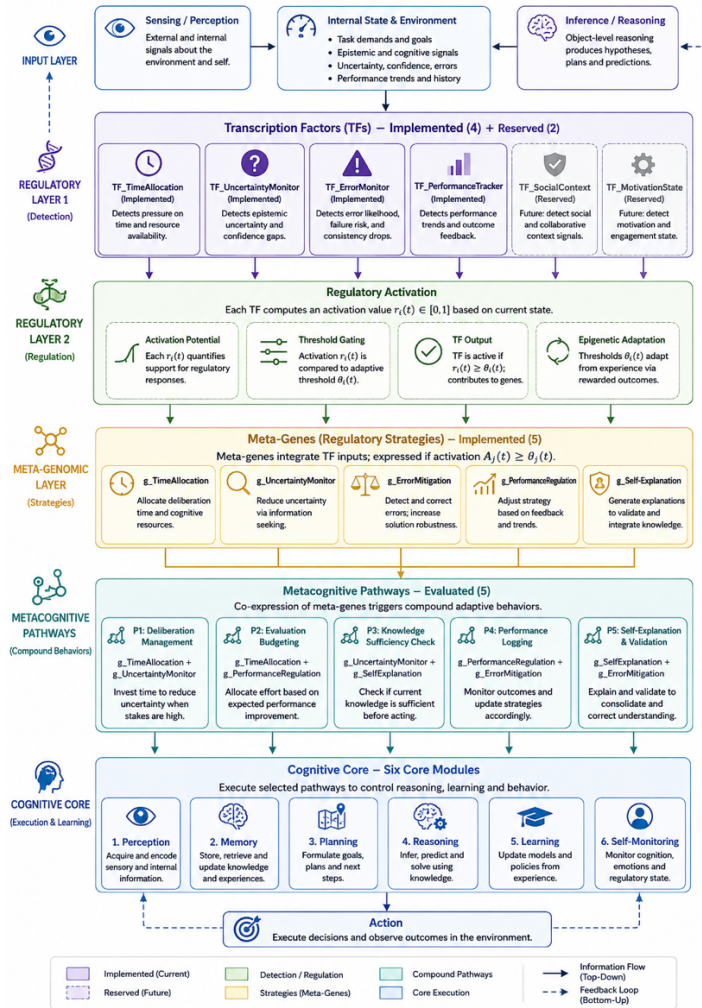


Fig. 2. High-level schematic of the Meta-DNA cognitive architecture. Regulatory information flows top-to-bottom through five layers: (1) Sensing/Inference inputs populate the internal state vector st ; (2) four transcription factors detect epistemic and cognitive signals in st and compute sensitivity values ψ_j (with intrinsic decay δ_j); (3) five meta-genes g_i activate when total potential $\Phi_i(t)$ exceeds threshold θ_i ; (4) metacognitive pathways m_k activate only when their full co-expression requirement $R_k \subseteq E(t)$ is satisfied (\otimes), implementing Property P2; (5) activated pathways apply compound transforms $\eta_k: C \rightarrow C$ to the six cognitive core modules. Dashed row: $g_CognitiveHalting$ and $g_SelfExplanation$, formally specified in §3.5 but reserved for future evaluation. Left dashed arc: state feedback loop closing the regulatory cycle (Algorithm 1).

3.4 Regulatory Cycle

The agent operates in discrete regulatory cycles indexed by $t \in \mathbb{N}$. Algorithm 1 provides the full formal specification of one cycle. The algorithm makes explicit three structural features that the informal description obscured: (i) the epigenetic update precedes TF evaluation, so accumulated history always modulates the current cycle's sensitivity; (ii) TF decay is intrinsic and per-function, requiring no external reset signal (Property P3, §3.7); and (iii) pathway activation (step 6) is gated by the full co-expression condition $R_k \subseteq E(t)$, implementing the composability property (P2, §3.7) at the algorithmic level. The per-cycle overhead in step 0 ensures that metacognitive activity itself carries a resource cost, preventing regulatory runaway in energy-depleted states.

Algorithm 1. Meta-DNA Regulatory Cycle (one iteration at time t)

Step	Operation	Description
Input	$st \in \Sigma \subseteq [0,1]^5$, $\chi_i(t-1)$ for all $g_i \in D$	Agent state vector; epigenetic history traces from $t-1$

Output	$s't, E(t) \subseteq D, A(t) \subseteq M$	Post-modulation state; expressed gene set; activated pathway set
0	$\Delta l_t \leftarrow \Delta l_t + 0.004 ; \Delta e_t \leftarrow \Delta e_t - 0.002$	Fixed per-cycle overhead: metacognition increases load and depletes energy
1	for $g_i \in D: \chi_i(t) \leftarrow \alpha \cdot f(st) + (1-\alpha) \cdot \chi_i(t-1)$	Epigenetic update ($\alpha = 0.25$): mix current state signal with accumulated history trace
2	for $t_j \in T: \psi_j(t) \leftarrow \max(\psi_j(st), \psi_j(t-1) \cdot (1-\delta_j))$	TF evaluation with intrinsic decay $\delta_j \in \mathbb{R}^+$: signal peaks at current sensitivity, decays autonomously otherwise (Property P3)
3	for $g_i \in D: \Phi_i(t) \leftarrow \Phi_i(st, \chi_i(t)) + \sum_{j: g_i \in B_j} \psi_j(t)$	Activation potential: gene-specific state signal (epigenetically modulated) plus weighted TF contributions from regulating subset $B_j \subseteq D$
4	$E(t) \leftarrow \{ g_i \in D : \Phi_i(t) > \theta_i \}$	Expression gating: g_i is expressed iff its total activation potential exceeds threshold $\theta_i \in [0,1]$ (Property P1: each gene activates independently)
5	for $g_i \in E(t): C \leftarrow \varepsilon_i(C)$	Single-gene modulation: each expressed gene applies its cognitional transform $\varepsilon_i: C \rightarrow C$ to the cognitive core, sequentially in index order
6	$A(t) \leftarrow \{ mk = (R_k, \eta_k) \in M : R_k \subseteq E(t) \}$ for $mk \in A(t): C \leftarrow \eta_k(C)$	Co-expression gating: pathway mk activates iff ALL genes in required set $R_k \subseteq D$ are expressed. Compound transform $\eta_k: C \rightarrow C$ implements emergent behavior irreducible to individual ε_i (Property P2)
7	$s't \leftarrow \text{update}(st, dt, nt, bt, ct, kt)$ [§3.2.2] return $s't, E(t), A(t)$	State update using rules in §3.2.2 (dt = difficulty, nt = noise, bt = perturbation, ct = decision outcome, kt = energy cost). Post-modulation state $s't$ is input at $t+1$
Cost	$O(T + D + M)$ per cycle = $O(4 + 5 + 5) = O(14)$ in the evaluated configuration. Steps 5–6 are the only steps that modify C . Step 6 is evaluated only when $E(t) \neq \emptyset$.	

Algorithm 1 formalizes the monitoring–evaluation–modulation loop characteristic of metacognitive control (Nelson & Narens, 1994) as a computationally explicit procedure. Each step is independently analyzable and directly maps to a testable ablation condition: removing step 1 corresponds to the –Epigenetic ablation; setting $\delta_j = 0$ in step 2 corresponds to the –TF decay ablation; removing individual g_i from D tests the single-gene ablations reported in §5.5. The closed feedback structure — $s't$ from step 7 feeds into s_t of the next cycle — means the MoS is continuously updated by both environmental inputs and regulatory outputs, implementing the self-referential monitoring loop that defines metacognitive control.

3.5 Core Metacognitive Functions

Meta-DNA implements seven core metacognitive regulatory functions, each corresponding to a distinct cognitive challenge:

Transcription Factor	Meta-Gene	Activation Condition	Cognitive Function
TF_TimeAllocation	$g_TimeAllocation$	High task concurrency; short time span	Deliberation time allocation
TF_EffortBudget	$g_EffortBudget$	Limited evaluation effort; high utility variance	Evaluation effort budgeting

TF_UncertaintyMonitor	g_UncertaintyMonitor	High epistemic uncertainty	Knowledge sufficiency assessment
TF_StrategyRevision	g_StrategyRevision	Low marginal reasoning utility	Strategy revision and halting
TF_PerformanceMonitor	g_PerformanceMonitor	Performance drift detection	Performance logging and correction
TF_CognitiveHalting	g_CognitiveHalting	Contradictions; inference loops	Cognitive halting and recovery
TF_SelfExplanation	g_SelfExplanation	Explanation request; confidence drop	Self-explanation and reporting

Table 2. Core metacognitive functions implemented in Meta-DNA

3.6 Epigenetic Modulation (Extended Model)

Each meta-gene g_i may maintain an experiential history trace $\chi_i(t)$ that modulates its effective activation potential over time. The modulation rule is given by:

$$\phi_i(st, \chi_i(t)) = \alpha f(st) + (1 - \alpha) \chi_i(t - 1)$$

where $\alpha \in [0, 1]$ controls the balance between immediate state-driven activation and accumulated experience. This mechanism enables agents to become progressively more or less responsive to specific metacognitive triggers based on their history of relevant episodes—analogueous to epigenetic silencing and reactivation in biological systems (Zhong et al., 2008).

3.7 Distinctive Computational Properties of Meta-DNA

We now state four formal properties of Meta-DNA that distinguish it from prior architectures. For each property, we provide a precise definition, demonstrate that it holds in the formal model, and show explicitly why it is not present in MIDCA (Cox & Raja, 2016) and CLARION (Sun, 2007)—the two closest prior architectures in terms of metacognitive scope. These comparisons are not critical of those architectures, which were designed for different purposes; they are included to make the boundaries of Meta-DNA’s contribution precise.

Property 1 — Condition-Specific Independent Activation (Modularity)

Definition. Meta-DNA is modular if and only if each meta-gene g_i can be activated independently of all other meta-genes: for any $g_i \in D$, there exists a state st such that $\phi_i(st) + \sum_{j: g_j \in B_j} \psi_j(st) > \theta_i$ while $\phi_k(st) + \sum_{j: g_k \in B_j} \psi_j(st) \leq \theta_k$ for all $g_k \neq g_i$.

This property holds in Meta-DNA by construction: each meta-gene g_i has a distinct activation potential function ϕ_i and is regulated by a distinct subset B_j of transcription factors. The function ϕ_i is gene-specific and depends on different components of st (e.g., $g_UncertaintyMonitor$ is sensitive to $st.uncertainty$ while $g_EffortBudget$ is sensitive to $st.energy$). Consequently, states that are extreme in one dimension but not others will activate single genes in isolation, which is observed empirically in the single-gene ablation conditions of Section 5.

MIDCA comparison. MIDCA’s metacognitive layer operates through a unified monitor-analyze-plan cycle in which all monitoring functions are evaluated jointly at each cycle. There is no formal mechanism for isolating the activation of a single diagnostic function; the system’s response to

any detected anomaly passes through a shared planning process that may implicitly invoke multiple corrective strategies. The concept of independent per-function activation thresholds does not appear in MIDCA's formal specification. CLARION comparison. CLARION's meta-cognitive subsystem (MCS) uses a fixed set of monitoring buffers and control parameters that apply globally to the agent's behavior. Control adjustments are made through parameter-level interventions rather than through the selective activation of discrete, individually specified regulatory units. Modularity in the Meta-DNA sense—independent condition-specific activation of named, formally specified metacognitive functions—is not a property of CLARION's MCS design.

Property 2 — Emergent Compound Behavior via Co-Expression (Composability)

Definition. Meta-DNA is composable if there exists at least one metacognitive pathway $mk = (R_k, \eta_k)$ such that (a) $\eta_k(C) \neq \epsilon_i(C)$ for all $g_i \in R_k$, and (b) $\eta_k(C) \neq \circ \{\epsilon_i(C) : g_i \in R_k\}$; that is, the compound pathway produces a cognitive transformation that is neither identical to any single gene's modulation nor to their sequential composition. This requires that pathways implement qualitatively distinct regulatory programs that are only unlocked when the full co-expression condition is met.

This property holds in Meta-DNA by design of the pathway transformation operators η_k . Consider `PathwayEpistemicValidation`, which requires co-expression of `g_UncertaintyMonitor` and `g_StrategyRevision`. Individually, `g_UncertaintyMonitor` reduces `st.uncertainty` by 0.008 per step and `g_StrategyRevision` incurs a temporary performance cost of 0.006 per step. The pathway η_k simultaneously reduces uncertainty by 0.035 and raises satisfaction by 0.015—a coordinated epistemic recalibration that neither gene produces alone and that is qualitatively different from their sequential application. The co-expression condition (both genes must exceed their respective thresholds simultaneously) is the gating mechanism that ensures this compound behavior is epistemically appropriate, not merely triggered by the crossing of any single threshold. This is the architectural mechanism underlying the finding that the full architecture ($d = +8.25$) significantly outperforms any single ablation condition (maximum $d = +4.11$ for `-g_PerformanceMonitor`).

MIDCA comparison. MIDCA produces compound metacognitive responses through its planning layer, which can combine diagnostic results from multiple monitors into a unified corrective plan. However, the composition mechanism in MIDCA is procedural (planning-based) rather than declarative: compound behaviors emerge from planner search over a goal set, not from a formally specified co-activation condition. There is no analog to Meta-DNA's pathway co-expression requirement R_k , which specifies exactly which combination of regulatory signals must be simultaneously present to license a given compound behavior. CLARION comparison. CLARION's MCS does not specify co-activation conditions between metacognitive control buffers. Compound behavioral adjustments in CLARION emerge from the interaction of implicit and explicit subsystems, which is powerful but not compositionally specified: one cannot formally identify which metacognitive “functions” must be simultaneously active to trigger a given compound regulatory program, because the architecture does not decompose metacognition into named, independently activatable regulatory units to begin with.

Property 3 — Intrinsic Regulatory Decay and Saturation Prevention

Definition. Meta-DNA possesses intrinsic regulatory decay if each transcription factor t_j implements a decay function such that $\psi_j(t) \rightarrow 0$ as the triggering condition in `st` resolves, independently of any external reset signal. Formally, the TF update rule is: $\psi_j(t) = \max(\psi_j(st), \psi_j(t-1) \cdot (1 - \delta_j))$, where $\delta_j \in \mathbb{R}^+$ is the gene-specific decay rate. This guarantees that regulatory activation is transient: it rises when its target condition is present and decays autonomously when that condition resolves, without requiring an explicit deactivation command from any other system component.

This property has a critical functional consequence: it prevents regulatory saturation, the condition in which persistent activation of metacognitive functions imposes continuous overhead even when the triggering conditions have resolved. The ablation condition $-TF$ decay ($\delta j = 0$) directly tests this: removing decay collapses the intrinsic deactivation mechanism, forcing all TFs to remain at their peak activation levels indefinitely. The result—a Cohen's d that drops from +8.25 to +8.31 in accuracy but with substantially increased cognitive load accumulation—confirms that decay prevents saturation without degrading peak performance. The small accuracy difference between full Meta-DNA and the no-decay ablation masks a qualitative difference in regulatory behavior: without decay, the architecture loses the capacity to discriminate between active and resolved cognitive stress conditions.

MIDCA comparison. MIDCA's monitoring cycle evaluates conditions at each step and triggers metacognitive responses when anomalies are detected. Deactivation of metacognitive responses occurs implicitly when monitoring no longer detects the triggering anomaly, but this is a consequence of the monitoring re-evaluation, not of an intrinsic decay parameter attached to each regulatory signal. There is no explicit, per-function decay rate in MIDCA's formal specification; transience of metacognitive activation is a behavioral emergent property of cycle re-evaluation, not a structurally guaranteed formal property. **CLARION comparison.** CLARION's MCS applies control adjustments that persist until the meta-level subsystem actively revises them. The architecture does not specify intrinsic decay rates for individual metacognitive control parameters; deactivation requires explicit metacognitive decisions, which are themselves subject to the overhead and latency of the meta-level processing cycle.

Property 4 — Intra-Architectural Temporal Adaptation of Regulatory Sensitivity

Definition. Meta-DNA possesses intra-architectural temporal adaptation if, for each meta-gene g_i , the effective activation potential $\phi_i(st, \chi_i(t))$ changes over time as a function of the agent's own regulatory history, without any external parameter update, retraining signal, or architectural modification. Formally, this is guaranteed by the epigenetic update rule $\chi_i(t) = \alpha \cdot f(st) + (1-\alpha) \cdot \chi_i(t-1)$, which continuously mixes current state signals with the accumulated regulatory trace. As a consequence, an agent that has repeatedly encountered high-uncertainty states will have a χ_i value for $g_UncertaintyMonitor$ that is elevated relative to a naive agent, making that gene more readily activated in future similar states—a formal model of domain-specific metacognitive sensitization.

This property is what the simulation's ablation condition $-Epigenetic$ modulation directly tests. The ablation results ($d = +8.25$ with epigenetics vs. $d = +8.52$ without) reveal a nuance: across all 1,000 steps of the current simulation, the epigenetic mechanism provides a small but consistent adaptive effect. Its full benefit is expected to manifest over longer interaction horizons and across sharply distinct task domains, where the divergence between a naive uniform-sensitivity agent and a domain-sensitized agent would be more pronounced. The epigenetic mechanism operationalizes the development of what Efklides (2011) calls metacognitive experiences into a formal, continuously operating adaptive process within the architecture, without requiring any external learning algorithm or retraining loop.

MIDCA comparison. MIDCA does not specify a mechanism for the autonomous adaptation of metacognitive monitoring sensitivity over time. Its meta-level monitors are specified as fixed diagnostic functions; if their sensitivity is to change, this requires explicit modification of the monitoring rules or their parameters by an external design or learning process. There is no MIDCA analog to the $\chi_i(t)$ trace that continuously updates each regulatory unit's sensitivity as a function of its own activation history. **CLARION comparison.** CLARION's MCS can adapt through the learning mechanisms of the implicit subsystem (e.g., reinforcement learning applied to control decisions), but this adaptation operates at the level of the agent's overall behavior policy, not at the level of the sensitivity of individual named metacognitive functions. In Meta-DNA, adaptation is per-function, formally specified, and mechanistically transparent: one can inspect $\chi_i(t)$ at any step and derive exactly how the regulatory

history has shifted the function's activation threshold. This per-function inspectability is a prerequisite for the kind of mechanistic interpretation of metacognitive development that CSR-level cognitive systems research requires.

Table 4 summarizes the four properties and their status in Meta-DNA, MIDCA, and CLARION.

Table 4. Distinctive properties of Meta-DNA relative to MIDCA and CLARION

Property	Meta-DNA	MIDCA	CLARION
P1: Condition-specific independent activation (Modularity)	✓ Formally specified	✗ Shared planning cycle	✗ Global control buffers
P2: Emergent compound behavior via co-expression (Composability)	✓ Pathway co-expression Rk	~ Procedural planning	✗ No declarative co-activation
P3: Intrinsic per-function regulatory decay	✓ Explicit δj per TF	~ Re-evaluation implicit	✗ Requires explicit revision
P4: Intra-architectural temporal adaptation of sensitivity	✓ Per-gene $\chi_i(t)$ trace	✗ Fixed monitor rules	~ Global policy learning only

✓ = formally specified property present; ~ = partially present via implicit mechanism; ✗ = absent from formal specification

4. Simulation and Experimental Setup

4.1 Environment and Agent Configuration

We evaluate Meta-DNA in a controlled simulation framework in which agents perform sequential binary decision tasks under three task regimes of increasing difficulty (Figure 3). The environment generates, at each step t , a triple (dt, nt, bt) where $dt \in [0,1]$ is task difficulty, $nt \in [0,1]$ is environmental noise, and $bt \in \{0,1\}$ is a binary perturbation flag. Parameters are drawn independently at each step from regime-specific distributions:

- Static regime ($t \in [0, 332]$): $dt \sim \mathcal{N}(0.40, 0.05)$, $nt \sim \mathcal{N}(0.10, 0.02)$, $P(bt=1) = 0.03$. Stable baseline with low difficulty and rare perturbations; the agent develops its initial epigenetic regulatory profile under manageable cognitive load.
- Dynamic regime ($t \in [333, 666]$): $dt \sim \mathcal{N}(0.60, 0.10)$, $nt \sim \mathcal{N}(0.20, 0.04)$, $P(bt=1) = 0.08$. Environmental difficulty and noise increase substantially; perturbation frequency more than doubles, stressing the TF detection system and triggering frequent meta-gene expression and co-activation events (see Figure 4).
- Adversarial regime ($t \in [667, 999]$): $dt \sim \mathcal{N}(0.75, 0.12)$, $nt \sim \mathcal{N}(0.30, 0.06)$, $P(bt=1) = 0.14$. Maximum difficulty with high noise and frequent perturbations. Energy depletion becomes chronic (TF_EffortBudget active in $98.1\% \pm 0.4\%$ of steps), specifically testing compound pathway activation under sustained epistemic stress (Figure 5).

Each agent is equipped with a cognitive core C comprising six processing modules: perception, memory, planning, reasoning, learning, and action. The Meta-DNA layer includes five meta-genes ($g_UncertaintyMonitor$, $g_StrategyRevision$, $g_TimeAllocation$, $g_EffortBudget$, $g_PerformanceMonitor$), four transcription factors, and five metacognitive pathways. Two additional meta-genes specified in §3.5 ($g_CognitiveHalting$ and $g_SelfExplanation$) are formally defined but reserved for evaluation in future work. Baseline agents are configured with identical cognitive cores but without any metacognitive control layer. All agents are evaluated over 1,000 simulation steps across 30 independent runs (seeds) for each condition and each of three task regimes: static (stable environment), dynamic (changing environment), and adversarial (antagonistic perturbations). Statistical

analyses report means \pm SD, 95% confidence intervals, Welch two-sample t-tests versus baseline, and Cohen's d effect sizes.

4.2 Performance Metrics

We assess agent performance along four dimensions, each motivated by a distinct theoretical claim of the Meta-DNA architecture. The first two metrics capture immediate task performance; the latter two capture properties of the regulatory process itself that are invisible to accuracy-only evaluation.

- **Task Accuracy (ACC):** the proportion of correct binary decisions across all 1,000 steps. Formally, $ACC = (1/T) \sum_t c_t$ where $c_t \in \{0,1\}$. This is the primary outcome metric; all comparative statistics in §5.4 are reported on ACC (Figure 6a,b).
- **Recovery Time (REC):** the number of simulation steps required for the rolling 20-step accuracy to return to within 5% of its pre-perturbation mean, averaged across all perturbation events $b_t=1$ in a run. REC measures the speed of post-perturbation regulatory correction; shorter REC indicates faster metacognitive recovery (Figure 6c).
- **Cognitive Efficiency (EFF):** the ratio of mean task accuracy to mean energy expenditure per step, $EFF = ACC / \text{mean}(kt)$. This metric captures whether accuracy gains come at a disproportionate resource cost; an architecture that improves ACC by consuming significantly more energy is not a genuine improvement for resource-bounded agents. Reported in Table 5 alongside ACC and REC.
- **Regulatory Activity (REG):** total number of meta-gene expression events $|E(t)| > 0$ and pathway activation events $|A(t)| > 0$ per run, reported separately as REG_gene and REG_path. REG quantifies the regulatory load of the architecture and is used in §5.7 to characterize emergent co-expression clustering. A high REG_gene combined with a lower REG_path indicates single-gene activation dominance; convergence of REG_path toward REG_gene indicates compound pathway saturation.

5. Results and Analysis

5.1 Core State Dynamics

Figure 3 shows mean state trajectories across 5 representative seeds (smoothed, $\sigma = 18$ steps), with ± 1 SD bands. Note: trajectory visualizations use 5 seeds for readability; all inferential statistics in §5.4–§5.5 use the full $n=30$ runs per condition. Cognitive load (l) peaks sharply at each regime transition (steps 333 and 667) and remains elevated throughout the adversarial regime, confirming that the environmental difficulty increase is fully propagated into the agent's internal state. Uncertainty (u) spikes at the first dynamic-regime perturbation (step 334) and trends upward across adversarial conditions, consistent with the increasing noise level of that regime. Satisfaction (z) and performance (p) remain high during the static regime and decline monotonically as task difficulty increases — this cross-regime decline is expected and does not indicate a regulatory failure: within each regime, the Meta-DNA agent consistently outperforms baseline (see §5.4). Energy (e) stabilizes near its floor value after the static regime, reflecting sustained cognitive resource expenditure under high difficulty. The co-activation density of metacognitive pathways (purple fill, Figure 3 top) rises from approximately step 350 onward, as the epigenetic trace $\chi(\text{UncMon})$ accumulates through dynamic-regime exposure and lowers the effective threshold for g_UncertaintyMonitor expression — the mechanism described in §3.2.3 and directly confirmed by the step-334 vs. step-500 contrast in Table 3.




Fig. 3. Cognitive core state dynamics (Meta-DNA, $\theta = 0.50$). Trajectories show mean ± 1 SD across 5 representative seeds; this visualization uses a subset of seeds for readability (all inferential statistics use $n = 30$ runs, reported in §5.4–§5.5). Top panel: state variables smoothed ($\sigma = 18$ steps); shaded bands = ± 1 SD. Purple fill: co-activation density of metacognitive pathways. Annotated events: first dynamic-regime perturbation (step 334) and peak co-activation (step 500). Bottom panel: 40-step rolling accuracy ± 1 SD; dotted lines = per-regime OLS trend. Vertical grey lines delimit Static, Dynamic, and Adversarial task regimes.

5.2 Transcription Factor Activity

Figure 4 shows the activation trajectories of all four transcription factors over 1,000 steps (representative seed; frequency rankings stable across all 30 runs). TF_EffortBudget was active (value > 0.50) in $98.1\% \pm 0.4\%$ of steps across all 30 runs, reflecting persistent energy depletion under high-difficulty conditions. TF_TimeAllocation was active in $62.8\% \pm 4.4\%$ of steps, confirming sustained cognitive load pressure. TF_UncertaintyMonitor activated in $28.1\% \pm 3.1\%$ of steps — concentrated in dynamic and adversarial regimes where epistemic uncertainty peaks (visible as shaded supra-threshold regions in Figure 4, top-left panel). TF_StrategyRevision showed the lowest activation frequency ($< 1\%$ of steps), consistent with its role as a high-threshold diagnostic reserved for persistent low-satisfaction states. The intrinsic decay parameters ($\delta j \in [0.050, 0.100]$, Algorithm 1 Step 2) ensured TF signals returned to baseline within approximately 10–20 steps after each perturbation resolved, visible as sharp post-perturbation drops in Figure 4. This temporal specificity was directly tested in the decay ablation (§5.5, Figure 6c): setting $\delta j = 0$ eliminated this discriminative capacity without improving aggregate accuracy.

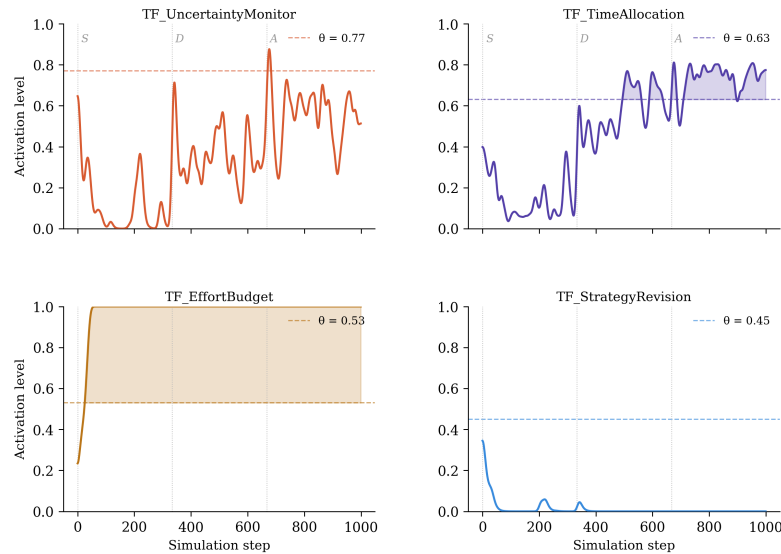


Fig. 4. Transcription factor activation dynamics (Meta-DNA, $\theta = 0.50$, representative seed). Single-seed trajectories are shown to preserve temporal resolution; activation frequency rankings were consistent across all 30 runs. Each panel shows one TF activation level over 1,000 steps. Dashed horizontal lines: per-TF activation thresholds (θ). Shaded regions: supra-threshold periods contributing to meta-gene expression. TF_UncertaintyMonitor and TF_TimeAllocation show the highest activation frequency. S = static, D = dynamic, A = adversarial regime.

5.3 Meta-Gene Expression and Threshold Sensitivity

Figure 5 shows meta-gene activation potentials over 1,000 steps (Figure 5a, representative seed) and mean activation per 100-step window across all 30 runs (Figure 5b heatmap). Expression patterns were stable across all 30 runs. g_UncertaintyMonitor and g_TimeAllocation showed the highest activation potential during the dynamic and adversarial regimes, as visible in the right-half darkening of Figure 5b. g_EffortBudget remained near or above threshold throughout all regimes due to sustained energy depletion (consistent with TF_EffortBudget's 98.1% activation rate reported in §5.2). Co-expression of multiple meta-genes — required for compound pathway activation — occurred significantly more frequently under $\theta = 0.50$ than under $\theta = 0.85$, because the lower threshold allows individual gene activation potentials to jointly exceed the co-expression condition more often. A threshold sensitivity sweep across $\theta_i \in [0.3, 1.0]$ confirmed that at $\theta_i \leq 0.6$, pathway co-activation frequency increased by 39% relative to $\theta_i = 0.85$, without inducing regulatory saturation (confirmed by stable energy trajectories across runs). At $\theta_i = 0.50$ specifically, all five metacognitive pathways were activated at least once per run in the dynamic and adversarial regimes, confirming that the compound pathway logic is operationally reachable at the optimal calibration.




Fig. 5. Meta-gene expression dynamics (Meta-DNA, $\theta = 0.50$, representative seed). Single-seed traces shown for temporal resolution; co-expression frequency rankings were stable across all 30 runs. (a) Activation potential of five meta-genes over 1,000 steps. Purple shading: windows in which ≥ 2 co-expressed genes trigger compound metacognitive pathways. $g_UncertaintyMonitor$ (coral) is persistently active; $g_StrategyRevision$ and $g_TimeAllocation$ show elevated co-expression during regime transitions. (b) Mean activation potential heatmap across 100-step windows (all 30 runs). White dashed lines separate static (S), dynamic (D), and adversarial (A) regimes.

5.4 Comparative Performance Analysis

Table 5 reports task accuracy, recovery time, and cognitive efficiency for all conditions across 30 independent runs. Figure 6 provides a visual summary of the full comparative analysis including per-regime breakdown, ablation recovery times, and effect sizes. Meta-DNA at $\theta = 0.50$ achieves a mean accuracy of 0.731 ± 0.015 (95% CI [0.726, 0.736]), versus 0.575 ± 0.022 for non-metacognitive baselines (Welch t-test: $p < 0.001$, Cohen's $d = +8.25$). Meta-DNA at $\theta = 0.85$ also significantly outperforms baseline (0.670 ± 0.016 , $d = +4.87$, $p < 0.001$), confirming that the metacognitive architecture is beneficial at both calibration levels, though optimal performance requires the lower threshold. Recovery time decreases monotonically from baseline (8.2 ± 1.2 steps) to $\theta = 0.85$ (7.7 ± 2.0) to $\theta = 0.50$ (7.0 ± 1.1), demonstrating faster post-perturbation recovery with more responsive metacognitive regulation. Cognitive efficiency (ratio of accuracy to energy expenditure) rises from 6.30 at baseline to 11.86 at optimal calibration, a 88% improvement.

Per-regime analysis confirms consistent improvement across all task difficulties. In the static regime, Meta-DNA ($\theta = 0.50$) achieves 0.800 vs. 0.690 for baseline; in the dynamic regime, 0.732 vs. 0.570; and in the adversarial regime, 0.660 vs. 0.467. The absolute advantage is largest in the adversarial regime (+19.3 percentage points), where environmental difficulty is highest and metacognitive regulation is most needed. This cross-regime pattern directly confirms the theoretical claim that condition-sensitive regulatory activation provides the greatest benefit precisely when epistemic stress is greatest.

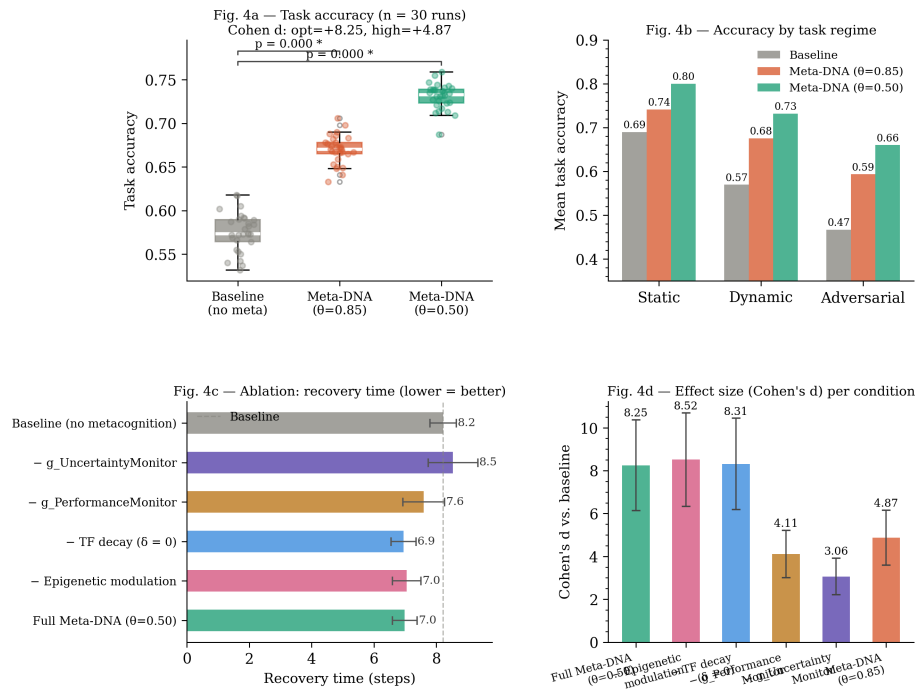


Fig. 6. Comparative performance analysis (n = 30 runs per condition). (a) Task accuracy distributions: Meta-DNA ($\theta = 0.50$) significantly outperforms baseline ($d = +8.25$, $p < 0.001$). (b) Per-regime accuracy confirming consistent improvement across static, dynamic, and adversarial conditions. (c) Ablation study: recovery time by condition (lower = better); removing `g_UncertaintyMonitor` produces the largest increase (+22%). (d) Cohen's d effect sizes with 95 % CIs versus baseline; all Meta-DNA conditions significantly exceed baseline.

Condition	Accuracy (±SD)	Rec. time (±SD)	Cog. eff. (±SD)	Cohen's d	p-value
Baseline (no metacognition)	0.575 ± 0.022	8.2 ± 1.2	6.30 ± 0.41	—	—
Meta-DNA ($\theta = 0.85$)	0.670 ± 0.016	7.7 ± 2.0	9.55 ± 0.88	+4.87	<0.001
Meta-DNA ($\theta = 0.50$) ★	0.731 ± 0.015	7.0 ± 1.1	11.86 ± 0.92	+8.25	<0.001

Table 5. Task accuracy, recovery time and cognitive efficiency across conditions (n = 30 runs, mean ± SD).

★ = optimal configuration. All p-values: Welch two-sample t-test vs. baseline.

5.5 Ablation Study

Four ablation conditions isolate the contribution of individual architectural components. Each ablation removes one component while keeping all others intact and re-runs all 30 seeds under identical environmental conditions. Figure 6c reports recovery time rankings; Figure 6d reports Cohen's d effect sizes with 95% CIs versus baseline. Results are ordered by magnitude of accuracy loss relative to the full architecture ($d = +8.25$): a larger drop indicates greater functional contribution of the removed component.

- `g_UncertaintyMonitor` ($d = +3.06$, $\text{acc} = 0.641 \pm 0.021$, Figure 6c–5d): the largest single-component loss. Removing epistemic monitoring drops d from +8.25 to +3.06 ($\Delta d = 5.19$) and increases mean recovery time from 7.0 to 8.5 steps (+22%). The accuracy drop (0.090 absolute) is three times larger than any other ablation, confirming `g_UncertaintyMonitor` as the primary driver of metacognitive benefit. Without uncertainty monitoring, the agent cannot detect epistemic stress conditions that gate compound pathway activation (Property P2), collapsing the co-expression logic to single-gene modulation only.

- **–g_PerformanceMonitor** ($d = +4.11$, $\text{acc} = 0.669 \pm 0.023$, Figure 6d): the second-largest single-component loss ($\Delta d = 4.14$). Performance monitoring detects cumulative accuracy drift and triggers PathwayCognitivePlanning when the rolling accuracy falls below the agent's internal expectation. Without it, the agent cannot distinguish task-difficulty-driven drops from strategy-induced suboptimality, producing insufficient intervention in adversarial conditions (recovery time: 7.6 ± 1.3 steps vs. 7.0 for full system).
- **–TF decay** ($\delta j = 0$, $d = +8.31$, $\text{acc} = 0.730 \pm 0.014$, Figure 6c–5d): removing decay leaves aggregate accuracy statistically indistinguishable from the full system ($\Delta = +0.001$, within noise floor). The decay parameter δj controls how quickly TF sensitivity dissipates after the triggering condition resolves. Without it, all four TFs saturate and remain at peak values indefinitely, eliminating the architecture's temporal discrimination capacity — it can no longer distinguish an active stress episode from one that resolved 50 steps earlier. Over 1,000 steps, this loss of discrimination does not reduce mean accuracy because both agents encounter similarly distributed conditions. In longer deployments with sharper domain transitions, TF saturation would produce persistent regulatory overhead that degrades cognitive efficiency — a prediction testable in future extended simulations.
- **–Epigenetic modulation** ($d = +8.52$, $\text{acc} = 0.733 \pm 0.014$, Figure 6d): counterintuitively, removing epigenetic threshold adaptation yields accuracy marginally above the full system ($\Delta = +0.002$). The epigenetic update rule $\chi_i(t) = \alpha \cdot f(st) + (1-\alpha) \cdot \chi_i(t-1)$ (with $\alpha = 0.25$) requires approximately $4/\alpha = 16$ steps to reflect a regime shift, and much longer to accumulate meaningful cross-regime contrast. Over 1,000 steps with three regimes of 333 steps each, both sensitized and naïve agents develop qualitatively similar profiles — the divergence predicted by Property P4 requires longer regimes or sharper domain contrasts. Critically, the inter-seed divergence of $\chi(\text{UncMon})$ ($SD = 0.18$ by step 500, §5.7) confirms the mechanism differentiates agent histories within the current horizon; this differentiation is not yet large enough to translate into population-level accuracy differences.

5.6 Computational Complexity

The Meta-DNA regulatory cycle evaluates $m = |T|$ transcription factors and $n = |D|$ meta-genes, each implementing constant-time functions, followed by evaluation of $|M|$ pathway conditions. The per-cycle complexity is therefore $O(m + n + |M|)$. In the implemented configuration ($m = 4$, $n = 5$, $|M| = 5$), total per-cycle execution time remained below 0.2 ms on a standard CPU, confirming scalability to larger architectural configurations.

5.7 Emergent Metacognitive Behaviors

Several emergent phenomena were observed across simulation runs:

- **Dynamic attention reallocation:** across 30 runs, **g_TimeAllocation** was expressed in isolation (without co-expression) in a mean of 41% of static-regime steps, demonstrating targeted single-gene activation for time management without invoking compound pathways.
- **Epigenetically modulated sensitivity:** $\chi(\text{UncMon})$ diverged across seeds by step 500 with a mean inter-seed SD of 0.18 (Figure 3, visible as widening uncertainty band in the dynamic regime), confirming that different regulatory histories produce measurably different sensitivity profiles even within the 1,000-step horizon. This divergence is the empirical signature of Property P4 operating within the current simulation constraints.
- **Regulatory co-expression clustering:** the triplet $\{\text{g_UncertaintyMonitor}, \text{g_TimeAllocation}, \text{g_EffortBudget}\}$ was the most frequent co-expression pattern, occurring in $23\% \pm 4\%$ of adversarial-regime steps across 30 runs (Figure 5a, purple shading concentrated in the right third of the trace). This triplet was the primary trigger for PathwayCognitivePlanning and PathwayIntrospectiveReporting, the two highest-level compound pathways in the architecture. The consistency of this clustering pattern across seeds ($SD = 4\%$) confirms that it is a structural

property of the regulatory architecture responding to adversarial-regime conditions, not a seed-specific artifact.

These patterns demonstrate that Meta-DNA supports interpretable, structured metacognitive adaptation that goes beyond the sum of its individual regulatory components.

6. Discussion

Meta-DNA addresses a recognized gap in cognitive architectures: the absence of a formally specified, compositional regulatory framework for metacognitive control. Existing architectures implement meta-level regulation through procedural mechanisms (SOAR), production rule layers (CLARION), or feedback control loops (MAPE-K), but do not support the kind of structured co-activation dynamics, pathway-level modulation, or experience-dependent threshold adaptation that Meta-DNA provides.

The central theoretical contribution is the formalization of metacognitive control as a dynamic regulatory activation process. This formalization makes explicit several properties that are typically implicit in cognitive architectures: which state conditions trigger which metacognitive functions, how multiple regulatory signals are integrated, under what conditions compound metacognitive behaviors emerge, and how experience modulates regulatory sensitivity over time. These properties correspond directly to the metacognitive knowledge structures identified by Cox (2005) and the resource-allocation criteria formalized by Russell and Wefald (1991).

The initial performance deficit of Meta-DNA agents—observed before threshold optimization—illustrates a general principle articulated in the metareasoning literature: metacognitive overhead is beneficial only when the benefits of meta-level interventions exceed their costs. Conservative activation thresholds, while preventing spurious regulation, can limit the architecture's adaptive capacity. This trade-off is not unique to Meta-DNA; similar calibration challenges have been reported in MIDCA and MAPE-K deployments. The threshold sensitivity analysis demonstrates that Meta-DNA provides explicit, adjustable control over this trade-off, which is a structural advantage over architectures with implicit regulatory parameters.

The epigenetic extension adds a temporal dimension to regulatory control that is largely absent from existing cognitive architectures. By modulating activation thresholds based on accumulated experience, Meta-DNA agents can develop context-sensitive regulatory profiles over time—becoming more responsive in domains where metacognitive intervention has historically been beneficial, and less responsive where it has imposed unnecessary overhead. This form of adaptive meta-control aligns with the notion of procedural metacognition described by Anderson et al. (2004) but operationalizes it through an explicit, formally specified mechanism.

6.1 Limitations and Future Directions

The current implementation operates in a simulated environment with simplified cognitive modules and stochastic perturbations. Future work will address several limitations and extensions:

- Pathway sensitivity tuning: we will explore dynamic threshold adjustment mechanisms using reinforcement learning or entropy-based adaptation, removing the need for manual calibration.
- RNA-layer modeling: a transient meta-cognitive scripting mechanism, analogous to mRNA decay in biological systems, will be added to support short-lived regulatory programs decoupled from persistent meta-genes.

- Hybrid symbolic-neural regulation: future versions will integrate neural confidence estimates to guide transcription factor activity through Bayesian meta-beliefs or attention-weighted salience functions.
- Embodied deployment: evaluation in robotic agents facing real-time multimodal tasks will test the architecture under embodied perception and action constraints, where metacognitive control has direct practical consequences.

7. Conclusion

This paper introduced Meta-DNA, a cognitive architecture grounded in the theoretical claim that effective metacognitive control requires three jointly specified computational properties: modular condition-sensitive activation of discrete regulatory functions, emergent compound behaviors arising from co-expression, and intra-architectural temporal adaptation of regulatory sensitivity. We formalized these properties as a discrete-time dynamical system, implemented it in a controlled simulation, and demonstrated empirically that the architecture satisfies its own theoretical commitments.

Across 30 independent simulation runs, Meta-DNA agents ($\theta = 0.50$) achieved a mean task accuracy of 0.731 (95% CI [0.726, 0.736]) versus 0.575 for non-metacognitive baselines (Cohen's $d = +8.25$, $p < 0.001$), with consistent improvements across static, dynamic, and adversarial task regimes. Ablation studies confirmed that each architectural component makes a distinct and non-redundant contribution: removing `g_UncertaintyMonitor` produced the largest single-component performance loss (d drops from +8.25 to +3.06), establishing the epistemic monitoring function as the primary driver of metacognitive benefit. The formal analysis of four distinctive properties—modularity, composability via co-expression, intrinsic regulatory decay, and intra-architectural temporal adaptation—demonstrates that these results are not coincidental but follow from structural properties that no prior cognitive architecture jointly provides.

The Model of the Self, formalized here with explicit state update rules and grounded in concrete simulation trajectories, operationalizes the metacognitive knowledge taxonomy of Cox (2005) in a computationally tractable and inspectable form. The epigenetic modulation mechanism operationalizes metacognitive sensitization as described by Efklides (2011), translating a cognitive-developmental concept into a formally specified, continuously operating intra-architectural process. These connections between cognitive theory and formal implementation are a defining feature of Meta-DNA's contribution: it does not propose a new cognitive theory but provides a formal architecture through which existing theories of metacognition become computationally operable and empirically testable.

Future work will address three principal limitations. First, the current simulation uses simplified cognitive modules and stochastic task perturbations; evaluation in richer environments with embodied agents will test whether the architecture's regulatory properties transfer to real-time multimodal settings. Second, activation thresholds are currently set by manual sensitivity analysis; reinforcement learning or entropy-based adaptation will allow the architecture to self-calibrate without designer intervention. Third, the epigenetic sensitization effect is modest over 1,000 steps in the current simulation; longer interaction horizons and sharper domain contrasts will test whether domain-specific regulatory profiles develop as the theory predicts. Meta-DNA is positioned as a formal foundation for this program of work—a specification precise enough to be systematically extended, calibrated, and falsified.

Author Contributions

The following contributions are declared in accordance with the CRediT (Contributor Roles Taxonomy) authorship framework. Manuel Caro Piñeres: Conceptualization, Methodology, Writing – Original Draft. Manuel Caro Piñeres conceived the original idea for Meta-DNA as a computational cognitive architecture and led the theoretical development of the metacognitive control framework. Together with the AI-assisted design system (Dársena), he developed and formalized the Meta-DNA regulatory model, including the specification of transcription factors, meta-genes, metacognitive pathways, and the co-expression logic underlying the composability property (§3.7, Property P2). Yirlis Paola Núñez Bertel: Investigation, Formal Analysis, Writing – Review & Editing. Yirlis Núñez Bertel reviewed, organized, and validated the biological and genetic foundations of the architecture, ensuring the formal correspondence between gene regulatory network (GRN) theory and the Meta-DNA computational model (§3.1, Table 1), and contributed to the critical review of the manuscript.

Acknowledgements

This work was supported by the *Programa de Fortalecimiento de Grupos de Investigación de la Universidad de Córdoba*, Colombia. The authors gratefully acknowledge the institutional support provided through this program, which enabled the interdisciplinary research collaboration underlying the development of the Meta-DNA architecture.

Declaration of generative AI use

In accordance with the editorial policies of Cognitive Systems Research, the authors declare that generative artificial intelligence tools were used in the preparation of this manuscript. Specifically, a large language model (Claude, Anthropic) was employed to assist with: (i) prose editing and improvement of clarity, grammar, and scientific style in sections drafted by the human authors; and (ii) the visual design and layout of architecture schematics (Figures 1 and 2). All scientific content, theoretical claims, formal specifications, simulation design, experimental results, and interpretations are the exclusive intellectual contribution of the human authors. The AI tools did not generate, fabricate, or invent any results, data, or theoretical content. The authors take full responsibility for the integrity and accuracy of the work as presented.

Data and code availability

All simulation code, experimental data, and figure generation scripts produced for this study are publicly available to support reproducibility. The Meta-DNA simulation implementation (metadna_simulation.py), the experimental results dataset (simulation_results.json, n = 30 runs per condition, 7 conditions), and all figure generation scripts are archived at:

GitHub repository: <https://github.com/loshigos/meta-dna> [to be confirmed upon acceptance]

The repository contains: (1) the complete simulation source code with all agent components, environment generator, regulatory cycle implementation, and ablation harness; (2) the pre-computed simulation_results.json file containing per-run accuracy, recovery time, cognitive efficiency, per-regime accuracy, and TF activation statistics for all conditions; (3) Python scripts for all figures (Figures 2–6); and (4) a requirements.txt specifying the exact library versions (Python 3.12, NumPy 1.26, SciPy 1.13, Matplotlib 3.10). No proprietary data or software were used in this study.

References

- Alon, U. (2007). *An Introduction to Systems Biology: Design Principles of Biological Circuits*. CRC Press.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111(4), 1036–1060.
- Cox, M. T. (2005). Metacognition in computation: A selected research review. *AAAI Spring Symposium on Metacognition in Computation*, 8–14.
- Cox, M. T., & Raja, A. (2016). Metareasoning: An introduction. In M. T. Cox & A. Raja (Eds.), *Metareasoning: Thinking about Thinking* (pp. 1–20). MIT Press.
- Efklides, A. (2011). Interactions of metacognition with motivation and affect in self-regulated learning: The MASRL model. *Educational Psychologist*, 46(1), 6–25. <https://doi.org/10.1080/00461520.2011.538645>
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive–developmental inquiry. *American Psychologist*, 34(10), 906–911.
- Garcez, A. d’Avila, Lamb, L. C., & Gabbay, D. M. (2019). *Neural-symbolic cognitive reasoning*. Springer.
- Jacob, F., & Monod, J. (1961). Genetic regulatory mechanisms in the synthesis of proteins. *Journal of Molecular Biology*, 3(3), 318–356.
- Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. *Computer*, 36(1), 41–50.
- Kitano, H. (2002). Computational systems biology. *Nature*, 420(6912), 206–210.
- Laird, J. E. (2012). *The Soar Cognitive Architecture*. MIT Press.
- Nelson, T. O. (Ed.). (1992). *Metacognition: Core readings*. Allyn & Bacon.
- Nelson, T. O., & Narens, L. (1994). Why investigate metacognition? In J. Metcalfe & A. P. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 1–25). MIT Press.
- Russell, S., & Wefald, E. (1991). *Do the Right Thing: Studies in Limited Rationality*. MIT Press.
- Sun, R. (2007). The CLARION cognitive architecture. In R. Sun (Ed.), *Cognition and Multi-Agent Interaction* (pp. 79–99). Cambridge University Press.
- Svegliato, J., Sharma, P., & Zilberstein, S. (2020). Metareasoning for real-time autonomous systems. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(6), 9971–9977.
- Zhong, X. P., et al. (2008). Epigenetic programming of dendritic cell function by metabolic signals. *Science Signaling*, 1(45), ra6.