

S-AI-ROBOTICS : A SPARSE ARTIFICIAL INTELLIGENCE ARCHITECTURE WITH HORMONAL ORCHESTRATION, PARSIMONIOUS CONTROL, AND SYMBOLIC MEMORY FOR ADAPTIVE, SAFE, AND EXPLAINABLE EMBODIED ROBOTICS

Said Slaoui

Mohammed V University, Rabat, Morocco

ABSTRACT

Robotic systems increasingly operate in dynamic, uncertain, and resource-constrained environments, where safety, energy efficiency, and explainability are as critical as raw performance. While learning-based and monolithic control architectures have demonstrated impressive capabilities, they often rely on continuous activation, data-intensive training, and opaque decision processes, making them fragile, energy-demanding, and difficult to audit in safety-critical contexts. This paper introduces S-AI-ROBOTICS, a bio-inspired and modular robotic intelligence framework grounded in the principles of Sparse Artificial Intelligence (S-AI). The proposed architecture departs from always-on robotic control by enforcing context-aware parsimony, where specialized robotic agents are activated only when justified by a symbolic hormonal state reflecting urgency, stability, energy, and confidence. A Robo-MetaAgent orchestrates agent activation through constrained optimization and hysteresis-based dynamics, ensuring stable and frugal behavior selection under competing objectives. To regulate behavioral priorities, S-AI-ROBOTICS integrates an artificial hormonal signaling layer, inspired by neuroendocrine systems, which modulates agent thresholds through bounded emission, inhibition, diffusion, and decay mechanisms. In parallel, a symbolic and contextual memory subsystem stores behavioral engrams—linking hormonal context, activated agents, actions, and outcomes—enabling rapid recall, adaptation, and native explainability of robotic decisions. The framework is evaluated using SAI-UT+, a reproducible experimental testbench, across multi-scenario robotic tasks including navigation, obstacle avoidance, energy scarcity, sensor degradation, and emergency handling. Results demonstrate that S-AI-ROBOTICS achieves improved stability, reduced energy consumption, and enhanced explainability compared to classical control, behavior trees, and reinforcement learning baselines, while maintaining robust performance under uncertainty. By unifying hormonal regulation, sparse orchestration, and symbolic memory within an embodied intelligence framework, S-AI-ROBOTICS establishes a principled foundation for adaptive, safe, and explainable robotic systems.

1. INTRODUCTION

1.1. From Reactive Control to Embodied Robotic Intelligence

Robotic systems have historically evolved from rigid, pre-programmed control schemes toward increasingly adaptive and autonomous architectures. Classical control paradigms—such as PID, state feedback, and model predictive control—offer strong guarantees of stability and precision, yet they remain fundamentally limited when operating in unstructured environments characterized by uncertainty, partial observability, and conflicting objectives. The emergence of cognitive robotics and learning-based approaches has significantly expanded robotic capabilities, enabling perception-driven decision-making, adaptive behaviors, and long-horizon planning.

However, this progress has come at a cost. Many modern robotic architectures rely on continuous activation, heavy data dependencies, and opaque internal representations, which challenge energy efficiency, safety certification, and explainability—especially in long-lived, embedded, or mission-critical robotic platforms. Robotic intelligence is inherently embodied: decisions are inseparable from physical constraints, sensor reliability, actuator limits, and energy availability. This embodiment demands not only adaptation, but also regulation—the ability to modulate which cognitive and behavioral processes are active, when, and to what extent.

1.2. Core Challenges in Modern Robotic Systems

Despite advances in learning and modular control, several persistent challenges remain unresolved:

1. **Context-Aware Parsimony** Most robotic systems remain computationally dense, activating perception, planning, and control modules continuously, even when the operational context does not justify their engagement.
2. **Behavioral Conflict and Instability** Robots must constantly arbitrate between competing objectives such as speed versus safety, exploration versus energy conservation, and reactivity versus stability. Without explicit regulation mechanisms, such conflicts often lead to oscillations or brittle behaviors.
3. **Energy and Resource Constraints** Autonomous robots—particularly mobile, edge, or field-deployed systems—operate under strict energy and computational budgets, making always-on intelligence unsustainable.
4. **Explainability and Trustworthiness** As robots increasingly interact with humans and operate in safety-critical environments, the inability to explain why a behavior was selected undermines trust, debugging, and certification.

These challenges point toward the need for robotic architectures that are not only adaptive, but also selective, regulated, and explainable by design.

1.3. Positioning of S-AI-ROBOTICS

This work proposes S-AI-ROBOTICS, an application of the Sparse Artificial Intelligence paradigm to embodied robotic systems. Rather than treating intelligence as a monolithic process, S-AI-ROBOTICS models robotic cognition as a sparse orchestration of specialized agents, each responsible for a distinct functional role such as navigation, obstacle avoidance, manipulation, safety monitoring, or energy management. At the core of the architecture lies a Robo-MetaAgent, which performs context-driven arbitration and activates only a minimal subset of agents required to address the current situation. This orchestration is guided by an artificial hormonal layer, inspired by biological endocrine systems, where symbolic hormonal variables encode macro-level robotic states such as urgency, fatigue, confidence, stability, and exploration drive. Hormones do not issue commands directly; instead, they modulate activation thresholds, priorities, and budgets, ensuring smooth, stable, and frugal behavioral transitions. Complementing this regulatory layer, S-AI-ROBOTICS integrates a symbolic and contextual memory that stores past behavioral episodes as structured engrams. These engrams link hormonal context, activated agents, executed actions, and observed outcomes, enabling rapid recall of validated strategies and providing intrinsic explainability for robotic decisions.

1.4. Contributions

The main contributions of this paper are:

1. A bio-inspired robotic intelligence architecture extending Sparse Artificial Intelligence to embodied robotic systems through hormonal regulation and sparse orchestration.
2. A formal hormonal signaling model for robotics, enabling adaptive modulation of behavioral priorities under uncertainty and resource constraints.
3. A parsimonious orchestration mechanism based on constrained optimization, hysteresis, and anti-flapping dynamics, ensuring stability and energy efficiency.
4. A symbolic memory framework for robotic behaviors, supporting engram-based recall, adaptation, and native explainability.

2. RELATED WORK

2.1. Classical and Hybrid Robotic Control Architectures

Robotic control has long relied on classical feedback and optimization-based paradigms, where stability, convergence, and constraint satisfaction are addressed through explicit mathematical formulations. Model Predictive Control (MPC) remains a cornerstone of this tradition, offering a principled framework for handling multivariable systems and constraints in real time [1], [2]. Recent advances in data-driven MPC extend these approaches to partially unknown or stochastic systems, improving adaptability while preserving formal guarantees [1]. Several surveys provide integrative perspectives on control strategies in robotics, covering linear, nonlinear, and hybrid formulations [3]. These works highlight the maturity and robustness of classical control pipelines, especially in structured environments. Foundational textbooks on MPC and robotic control further formalize these methods, emphasizing separation between modeling, control, and execution layers [4], [5]. Despite their rigor, classical feedback systems are inherently always-on, with limited capacity for contextual modulation or selective activation. Hybrid control approaches attempt to combine discrete logic with continuous control laws, yet they often rely on predefined switching rules and static hierarchies [6], [7]. As a result, such architectures struggle to scale toward complex autonomous behavior under uncertainty, energy constraints, and conflicting objectives. Nonlinear MPC frameworks improve expressiveness but remain computationally demanding and lack intrinsic mechanisms for behavioral explainability or adaptive parsimony [8].

2.2. Learning-Based And Reinforcement Learning Robotics

Learning-based control, and reinforcement learning (RL) in particular, has become a dominant paradigm for robotic autonomy, enabling agents to acquire control policies directly from interaction data. Comprehensive surveys document the breadth of RL algorithms and their applications to robotic navigation, manipulation, and control [9], [10]. Deep reinforcement learning (DRL) further expands representational capacity, allowing end-to-end learning from high-dimensional sensory inputs [10]. More recent surveys emphasize explainability and interpretability issues in RL, highlighting the opacity of learned policies and the difficulty of attributing decisions to internal representations [11]. In autonomous systems, RL has been applied to optimal decision-making under uncertainty, with promising results in simulation-dominated settings [12]. However, earlier foundational works already identified core challenges related to sample inefficiency, safety during exploration, and poor generalization outside training distributions [13]. Safety-aware learning has therefore emerged as a critical research direction, introducing constraints, shields, and conservative updates to mitigate catastrophic failures [14], [15]. Nevertheless, most safe RL approaches still rely on dense policy execution and require extensive retraining when system conditions change. Canonical DRL algorithms, such as entropy-regularized actor-critic methods, remain computationally intensive and difficult to

certify [16], [17]. Even successful demonstrations of end-to-end visuomotor policies reveal fragility under distribution shifts and limited explainability at the behavioral level [18].

2.3. Modular, Behavior-Based, and Subsumption Architectures

In parallel with learning-based approaches, modular and behavior-based architectures have sought to decompose robotic intelligence into interacting functional units. Behavior Trees (BTs) have gained popularity as a structured and human-readable alternative for sequencing and arbitration of robot behaviors [19]. Empirical evaluations demonstrate their practical advantages in terms of modularity and debugging [20]. Formal treatments of BTs clarify their execution semantics and compositional properties, positioning them as an intermediate abstraction between finite state machines and planning systems [21]. Earlier behavior-based paradigms, such as the subsumption architecture, introduced layered control structures where lower-level reflexes could override higher-level behaviors [22]. These ideas strongly influenced modern robotics, particularly in reactive and embodied systems.

Hierarchical control and whole-body behavior synthesis further extend modular architectures by coordinating multiple control objectives through prioritized task formulations [23]. Planning-based frameworks integrate symbolic reasoning and motion generation, offering expressive task representations [24], [25]. However, as architectural complexity increases, modular systems often rely on static priorities, handcrafted arbitration rules, or global planners that are continuously active. Reference architectures for distributed manufacturing and holonic systems illustrate these limitations at scale, where coordination overhead and rigidity become dominant concerns [26].

2.4. Bio-Inspired and Neuroendocrine Robotic Systems

Bio-inspired robotics explores mechanisms derived from biological systems to enhance adaptability and robustness. Neuroendocrine models introduce artificial hormonal signals as global modulators of behavior, influencing perception, decision-making, and action selection [27]. Such approaches aim to capture long-term regulation, motivation, and homeostasis beyond purely reactive control. Recent work on motivational and biologically inspired models demonstrates the potential of artificial hormones to regulate autonomous behavior in dynamic environments [28]. Extensions toward social and human-robot interaction contexts further highlight the role of neuroendocrine mechanisms in mediating affective and cooperative behaviors [29]. Foundational perspectives on embodied intelligence emphasize that cognition emerges from the tight coupling between body, environment, and control dynamics [30]. Evolutionary and adaptive robotics research complements these ideas by exploring population-based and developmental mechanisms for behavioral adaptation [31]. Within this landscape, Sparse Artificial Intelligence (S-AI) introduces a principled framework for parsimonious decision-making orchestrated by a hormonal MetaAgent, originally formulated independently of robotic embodiment [32]. Subsequent work formalizes hormonal modulation and adaptive orchestration mechanisms in conversational and symbolic settings [33], while domain-specific applications demonstrate scalability in distributed networking environments [34]. However, these contributions stop short of addressing the full spectrum of embodied robotic constraints.

2.5. Explainable, Safe, and Energy-Aware Robotics

As robotic systems increasingly operate in human-centered and safety-critical environments, explainability, safety, and energy efficiency have become first-class requirements. Recent manifestos on Explainable Artificial Intelligence (XAI) argue for intrinsic, rather than post-hoc, explanations grounded in system dynamics and decision processes [35]. Surveys dedicated to

explainable robotics confirm that most existing methods rely on external explanation layers, disconnected from the control architecture itself [36]. Causal approaches to explaining robot failures provide deeper insights into decision breakdowns, yet they are typically applied after execution rather than during control [37]. Broader surveys on safe and explainable robotics highlight the fragmentation of current solutions, where safety constraints, interpretability, and performance are often treated independently [38]. From a control perspective, safety-aware learning and control frameworks emphasize constraint satisfaction and risk reduction, but frequently at the cost of computational overhead and reduced flexibility [39]. Energy-aware robotics research addresses efficiency at the software and control levels, proposing optimization techniques and architectural guidelines for reducing power consumption [40]. Despite these advances, a unified framework integrating safety, explainability, and parsimony into the core decision process remains largely absent.

2.6. Identified Gaps and Positioning

The reviewed literature reveals a persistent gap between expressive control or learning capabilities and the requirements of adaptive, safe, and explainable embodied intelligence. Classical control offers rigor but lacks contextual modulation, learning-based approaches provide flexibility at the expense of safety and interpretability, and modular architectures struggle with scalability and global arbitration. Bio-inspired systems introduce promising regulatory mechanisms, yet often remain fragmented or weakly grounded in symbolic reasoning. Finally, explainability and energy efficiency are typically addressed as external constraints rather than intrinsic properties of the architecture. These gaps motivate the need for a unified, parsimonious, and hormonally regulated framework for embodied robotic intelligence.

3. THEORETICAL AND BIO-INSPIRED FOUNDATIONS

3.1. Parsimony as a Structuring Principle of Embodied Intelligence

In embodied robotic systems, parsimony is not a mere computational optimization but a fundamental organizational principle. An autonomous robot does not operate in an abstract, unlimited space; it is subject to strict physical, energetic, temporal, and material constraints [5]. Any unnecessary activation of perceptual, decisional, or motor processes directly translates into energy waste, component wear, or behavioral instability [40]. Most classical robotic architectures—whether control-based or learning-driven—adopt a dense paradigm, in which all cognitive modules remain continuously active regardless of contextual relevance [3], [13]. While straightforward to implement, this approach becomes increasingly inefficient and brittle in dynamic and uncertain environments. In S-AI-ROBOTICS, parsimony is elevated to a structural property of intelligence: a robotic module should be active only if it provides a measurable decisional benefit in the current context [32]. This principle leads to a selective orchestration of robotic capabilities, where intelligence emerges not from computational abundance, but from contextually justified activation. Parsimony thus becomes a necessary condition for robustness, energy sustainability, and explainability [35].

3.2. Cognitive Modularity and Specialized Robotic Agents

Robotic intelligence cannot be efficiently designed as a monolithic process. Robotic tasks, perception, navigation, manipulation, safety monitoring, and energy management, exhibit fundamentally different dynamics, time scales, and criticality levels [5]. S-AI-ROBOTICS adopts a strict cognitive modularity, in which each robotic capability is encapsulated within a specialized autonomous agent [19]. Each agent is characterized by:

- a clearly delimited functional role,
- a restricted set of perceptual inputs,
- a well-defined action space,
- explicit activation rules.

Unlike classical multi-agent robotic architectures, where coordination relies on explicit negotiation or communication protocols, S-AI-ROBOTICS introduces indirect coordination mediated by a shared hormonal state [27]. Agents do not coordinate with one another directly; instead, they are co-regulated through global hormonal variables. This design choice drastically reduces interaction complexity and mitigates decisional conflicts. Such regulated modularity enables:

- selective activation,
- improved fault isolation,
- a clear traceability between context, decision, and action [36].

3.3. Artificial Hormonal Signaling for Behavioral Regulation

Inspired by biological neuroendocrine systems, artificial hormonal signaling constitutes the central regulatory mechanism of S-AI-ROBOTICS [27]. Unlike direct neural-like signals, hormones do not prescribe actions; they modulate the conditions under which actions become possible. In the robotic context, an artificial hormone is defined as a bounded symbolic variable representing a macro-behavioral state of the system [28]. Hormonal variables evolve continuously or discretely through:

- local emissions derived from sensors and diagnostics,
- temporal decay mechanisms,
- cross-inhibition processes,
- saturation effects preventing excessive reactions.

Hormones never act directly on actuators. Instead, they: • adjust agent activation thresholds, • modulate relative priorities, • constrain energy and decisional budgets [40].

This regulatory scheme provides three major advantages:

1. Behavioral stability, through smooth and gradual transitions.
2. Robustness to uncertainty, by filtering noisy fluctuations.
3. Native explainability, as hormonal variables remain interpretable by design [35].

3.4. Embodied Intelligence and Body–Environment–Decision Coupling

Robotic intelligence is intrinsically embodied: every decision is conditioned by the robot's physical state and its interaction with the environment [30]. Cognitive architectures that ignore this embodiment inevitably produce unrealistic or unstable behaviors.

S-AI-ROBOTICS explicitly integrates this coupling through:

- Robotic Gland Agents, which translate bodily and environmental signals (fatigue, instability, danger, perceptual uncertainty) into hormonal emissions,
- a continuous feedback loop between perception, hormonal regulation, and orchestration [27].

As a result, the robot's internal state (energy level, sensor reliability, mechanical stress) directly influences cognitive decisions without relying on ad hoc heuristic layers. This enables behaviors such as:

- spontaneous speed reduction under instability,
- prioritization of safety over performance,
- controlled reduction of exploration under energy scarcity [39].

Embodied intelligence thus becomes self-regulated, rather than merely reactive.

3.5. Regulation, Stability, and Prevention of Behavioral Oscillations

A central challenge in adaptive robotic systems is the prevention of behavioral oscillations, such as rapid switching between contradictory strategies, unstable priority shifts, or overreactions to minor perturbations [7].

S-AI-ROBOTICS addresses this challenge structurally by combining:

- bounded hormonal variables,
- hysteresis mechanisms in agent activation,
- minimum activation durations (time-to-live),
- cooldown periods following deactivation.

These mechanisms ensure that robotic decisions evolve on time scales compatible with the physical dynamics of the robot and its environment [6]. Stability is no longer delegated solely to low-level control loops, but emerges as a property of cognitive orchestration itself [33].

4. GLOBAL ARCHITECTURE OF S-AI-ROBOTICS

4.1. System Overview and Architectural Layers

S-AI-ROBOTICS is structured as a layered, bio-inspired architecture designed to support regulated, parsimonious, and explainable robotic intelligence [27], [32]. The architecture is composed of the following layers:

1. Perception and Probing Layer
2. Robotic Gland Layer
3. Hormonal Engine (Graph-Based)
4. Parsimonious Orchestration Layer (Robo-MetaAgent)
5. Specialized Robotic Agents Layer
6. Symbolic and Contextual Memory Layer

Each layer operates at a distinct temporal and functional scale. Intelligence emerges from regulated interactions between layers rather than from continuous activation of all components [3].

4.2. Robo-Meta Agent: Parsimonious Orchestration Core

The Robo-MetaAgent is responsible for selecting and regulating specialized robotic agents under hormonal, energetic, and safety constraints [23]. Let $A=\{a_1,a_2,\dots,a_N\}$ denote the set of robotic agents..

At decision cyclet, the Robo- Meta Agent selects an activation vector:

$$\mathbf{x}(t)=(x_1(t),\dots,x_N(t)), \quad x_i(t)\in[0,1]$$

The orchestration problem is formulated as:

$$\max_{x(t)} \sum_{i=1}^N x_i(t)u_i(h(t),s(t)) \text{ subject to: } \sum_{i=1}^N x_i(t)c_i \leq b(t)$$

where u_i denotes the utility of agent ai , c_i its cost, and $B(t)$ the available energy and safety budget.

A continuous primal–dual dynamic governs convergence: $\dot{x}_i = \eta_i \Pi[0,1](u_i - \lambda c_i - \rho a x_i)$

$$\dot{\lambda} = \eta_\lambda (\sum_i x_i c_i - B(t))_+$$

ensuring sparse and stable agent activation [6].

4.3. Robotic Gland Agents and Hormonal Emission

S-AI-ROBOTICS introduces five canonical robotic hormones:

- Taskin: task demand and urgency
- Hazardin: hazard and collision risk
- Energexin: energy and thermal budget
- Resiliencin: stabilization and recovery
- Normin: mission and safety compliance

Robotic Gland Agents monitor local observables (tracking error, obstacle proximity, force/torque spikes, battery level, temperature, rule violations) and emit hormones via bounded logistic functions [27], [28]. For gland gj and hormone hk , local emission at node i is defined as:

$$e_{j,k,i}(t) = \sigma(\alpha_{j,k} z_{j,i}(t) - \beta_{j,k})(1 - h_{k,i}(t))$$

where $z_{j,i}(t)$ aggregates normalized robotic observables and the saturation term prevents runaway activation.

4.4. Hormonal Engine: Graph-Based Reaction–Diffusion Dynamics

Hormonal regulation is modeled as a reaction–diffusion process over a robotic graph with nodes i [27]. For each hormone hk at node i :

$$\dot{h}_{k,i}(t) = e_{k,i}(t) - \lambda_k h_{k,i}(t) - \sum_{\ell \neq k} \gamma_{k\ell} h_{\ell,i}(t) + D_k \sum_j L_{ij} h_{k,j}(t) - \chi_i(t) + \xi_{k,i}(t)$$

where: - λ_k is the decay rate, - $\gamma_{k\ell}$ models cross-inhibition, - L is the normalized graph Laplacian, - $\chi_i(t)$ encodes local physical constraints (battery, thermal, compute, actuator saturation), - $\xi_{k,i}(t)$ is bounded noise.

4.5. Discrete-Time Scheme and Stability

For SAI-UT+ simulations, dynamics are discretized using an Euler–Maruyama scheme with projection: $h_{k,i}(t+1)=\Pi_{[0,1]}(h_{k,i}(t)+\Delta th_{k,i}(t)+\sigma_k\sqrt{\Delta t} \varepsilon_{k,i}(t))$

Stability is ensured under standard explicit Euler conditions on the step size Δt and normalized Laplacian bounds [8].

4.6. Specialized Robotic Agents and Hysteresis

Each robotic agent a_i is activated based on hormonal thresholds with hysteresis:

$$a_i(t^+) = \begin{cases} 1 & \text{if } x_i(t) \geq \theta_i^\uparrow \\ 0 & \text{if } x_i(t) \leq \theta_i^\downarrow \\ a_i(t^-) & \text{otherwise} \end{cases}$$

with $\theta_i^\uparrow > \theta_i^\downarrow$, preventing oscillatory behavior [7].

4.7. Symbolic and Contextual Memory

Each decision cycle produces a symbolic engram: $\mathcal{E}m = \langle h(t), \mathbf{x}(t), a(t), o(t) \rangle$

Engrams are stored in a Dynamic Contextual Memory and retrieved via similarity metrics to bias future orchestration while preserving explainability [35], [36].

4.8. End-to-End Decision Cycle

The complete S-AI-ROBOTICS cycle consists of:

1. Perception and probing
2. Hormonal emission
3. Graph-based hormonal update
4. Parsimonious orchestration
5. Agent activation and action execution
6. Memory encoding and feedback

This architecture enforces regulated, frugal, and stable robotic behavior by construction [32].

5. TYPOLOGY OF SPECIALIZED ROBOTIC AGENTS

This section details the library of specialized agents $A = \{a_1, \dots, a_N\}$ employed by S-AI-ROBOTICS. In accordance with Section 4, each agent is (i) modular, (ii) parsimoniously activable via the activation vector $\mathbf{x}(t) \in [0,1]^N$, (iii) hormonally regulated through the robotic hormones {Taskin, Hazardin, Energexin, Resiliencin, Normin}, and (iv) stabilized by hysteresis at the execution level through thresholds $\theta_i^\uparrow, \theta_i^\downarrow$ [19],[32].

Each agent a_i is characterized by the following 6-tuple: $a_i \equiv \langle \mathcal{I}_i, \mathcal{O}_i, \mathcal{U}_i, \mathcal{C}_i, \Theta_i, \Pi_i \rangle$

where \mathcal{I}_i denotes the set of inputs (observables/percepts), \mathcal{O}_i the set of outputs (actions/commands/intentions), \mathcal{U}_i the utility function $u_i(h(t), s(t))$, \mathcal{C}_i the activation cost c_i , Θ_i the threshold and hysteresis parameters, and Π_i a set of local policies (control, planning, symbolic rules).

As defined in Section 4, orchestration selects $\mathbf{x}(t)$ via constrained optimization:

$$\max_{\mathbf{x}(t)} \sum_{i=1}^N x_i(t) u_i(h(t), s(t)) \quad \text{s.t.} \quad \sum_{i=1}^N x_i(t) c_i \leq B(t)$$

and the final ON/OFF decision is stabilized by hysteresis:

$$a_i(t^+) = \begin{cases} 1 & \text{if } x_i(t) \geq \theta_i^\uparrow \\ 0 & \text{if } x_i(t) \leq \theta_i^\downarrow \\ a_i(t^-) & \text{otherwise} \end{cases}$$

The typology introduced below is therefore explicitly designed to be compatible with $u_i(\cdot)$, c_i , and Θ_i , while remaining consistent with the graph-based hormonal engine (Sections 4.4–4.5).

5.1. Design Principles of Robotic Agents

5.1.1. Agents as Action Modules under Endocrine Regulation

Each agent implements an elementary or composite robotic capability (navigation, avoidance, manipulation, etc.) but must not be continuously active. Availability is permanent; activation is sparse [27], [32].

Two complementary mechanisms enforce this parsimony: 1. Hormonal pressure (pre-activation): hormones modulate $u_i(h(t), s(t))$ and/or thresholds Θ_i . 2. Budget and cost constraints (effective activation): c_i and $B(t)$ restrict the selected agent set.

5.1.2. Explicit Coupling with Robotic Hormones

Let $h(t)$ denote the global hormonal state (aggregated), and $hk, i(t)$ the node-level hormonal state defined in Section 4.4. For orchestration, a robust aggregation is used: $h^- k(t) = \text{Agg}(\{hk, i(t)\})$ where $\text{Agg}(\cdot)$ may be a trimmed mean, quantile, or bounded maximum, depending on the required conservatism.

The agents are defined to respond to the following hormones: - Taskin: task urgency, backlog, tracking error. - Hazardin: obstacle proximity, collision probability, force–torque spikes. - Energexin: energy, thermal, and power constraints. - Resiliencin: stabilization and recovery demand. - Normin: mission, safety, and HRI compliance [28].

5.2. Agent Registry and Interface Contracts

5.2.1. Canonical Registry

A minimal yet expressive registry A is structured into the following families:

- Task and planning agents: $\{a^{\text{task}}\}$
- Safety and avoidance agents: $\{a^{\text{safety}}\}$
- Locomotion and navigation agents: $\{a^{\text{nav}}\}$
- Manipulation and interaction agents: $\{a^{\text{manip}}\}$
- Perception and estimation agents: $\{a^{\text{percep}}\}$
- Resilience and recovery agents: $\{a^{\text{res}}\}$
- Energy and thermal agents: $\{a^{\text{energy}}\}$
- Compliance and governance agents: $\{a^{\text{norm}}\}$

- Super vision and local arbitration agents: $\{a^{\text{sup}}\}$ [19].

5.2.2. Input–Output Contract

Each agent consumes a subset $si(t)$ of global observables $s(t)$ and produces an output $yi(t)$ in the form of symbolic intentions and/or continuous commands.

Formally: $yi(t)=\pi i(si(t),h^-(t),m(t))$

where $m(t)$ is the contextual memory state (DCM) and πi the internal policy of agent ai .

The Robo-MetaAgent then aggregates: $u(t)=\text{Fuse}(\{ai(t) yi(t)\}i)$

where $\text{Fuse}(\cdot)$ is a robust fusion operator (priority rules, safety-max, hierarchical composition) [23]..

5.3. Detailed Specification of Agent Families

For each agent ai , we specify: - Role and objectives. - Key inputs $si(t)$. - Outputs $yi(t)$. - Dominant hormones and effects. - Cost ci . - Policy πi . - Interaction with hysteresis and stability mechanisms.

5.3.1. Task Management Agent: TaskManagerAgent

Role. Maintain mission execution (goals, backlog, deadlines) and produce high-level intentions (waypoints, subtasks, priorities).

Inputs. Tracking error $etrk(t)$, backlog $b(t)$, estimated delays $d(t)$, mission progress.

Outputs. Mission intention $ytask(t)$.

Hormonal coupling. - Taskin increases $utask$ with rising $etrk$, b , or d . - Normin constrains admissible actions. - Hazardin inhibits aggressive modes.

Utility (example).

$$u_{\text{task}}(t)=w_T \bar{h}_{\text{Taskin}}(t)-w_H \bar{h}_{\text{Hazardin}}(t)-w_N \bar{h}_{\text{Normin}}(t)$$

Cost. Moderate to high (planning and replanning).

Stability. Strong hysteresis when safety-critical regimes dominate.

5.3.2. Navigation Agent: Navigation Agent

Role. Generate feasible trajectories and nominal control under dynamic constraints.

Inputs. Local map, pose, velocity, obstacles, waypoints, uncertainty estimates.

Outputs. Navigation command $ynav(t)$.

Hormonal coupling. - Taskin increases urgency and directness. - Energexin constrains speed and favors energy-efficient paths. - Normin enforces regulatory constraints.

Cost. Medium.

Policy. Robust trajectory following with local planning [2].

5.3.3. Obstacle Avoidance Agent: ObstacleAvoidanceAgent

Role. Guarantee short-term collision avoidance and override nominal motion if necessary.

Inputs. Obstacle proximity $p_{obs}(t)$, collision probability $P_{col}(t)$, range sensors, velocity.

Outputs. Safety command $y_{avoid}(t)$.

Hormonal coupling. - Hazardin strongly increases u_{avoid} . - Taskin is inhibited by Hazardin. - Normin reinforces conservative behavior.

Cost. Low to medium, but high priority.

Stability. Tight hysteresis to avoid oscillations near sensing limits [7].

5.3.4. Compliance Agent: ComplianceAgent

Role. Enforce mission rules, safety regulations, and HRI constraints.

Inputs. Rule set R , robot state, violation indicators $v(t)$.

Outputs. Constraint masks $y_{norm}(t)$.

Hormonal coupling. - Normin dominates activation. - Hazardin may strengthen conservative constraints.

Cost. Low.

Interaction. Acts as a hard constraint during fusion.

5.3.5. Energy and Thermal Management Agent: EnergyThermalAgent

Role. Monitor and regulate energy and thermal budgets.

Inputs. Battery level $batt(t)$, temperature $T(t)$, power draw $P(t)$, compute load.

Outputs. Energy constraints $y_{energy}(t)$.

Hormonal coupling. - Energexin increases under low battery or overheating. - Energexin reduces $B(t)$ and contributes to $\chi_i(t)$ in the hormonal engine.

Cost. Low to medium.

5.3.6. Recovery Agent: RecoveryAgent

Role. Stabilization and recovery after failure, near-failure, or instability.

Inputs. Failure flags, instability metrics, degraded perception, near-collision events.

Outputs. Recovery plan $y_{rec}(t)$.

Hormonal coupling. - Resiliencin increases upon instability or recovery success. - Hazardin may trigger immediate activation. - Taskin is temporarily inhibited.

Cost. Medium to high.

Stability. Wide hysteresis and minimum activation duration.

5.3.7. Perception Health Agent: Perception Health Agent

Role. Assess sensor reliability and detect degradation.

Inputs. SNR, dropout rate, multi-sensor inconsistencies.

Outputs. Reliability scores and masks $y_{percep}(t)$.

Hormonal coupling. - Hazardin increases under uncertainty near obstacles. - Resiliencin increases under critical sensor failure. - Normin reinforces conservative behavior in HRI contexts.

Cost. Low to medium.

5.3.8. Manipulation Agent: ManipulationAgent

Role. Control grasping and end-effector trajectories.

Inputs. Target pose, force–torque signals, contact constraints, vision.

Outputs. Manipulation command $y_{manip}(t)$.

Hormonal coupling. - Taskin drives execution priority. - Hazardin increases with unexpected contacts. - Normin enforces HRI and safety limits. - Energexin reduces aggressiveness under constraints.

Cost. Medium to high.

5.3.9. Human Interaction Agent: HumanInteractionAgent

Role. Regulate human–robot interaction (distance, speed, social compliance).

Inputs. Human detection, proxemics, gestures, safety zones.

Outputs. HRI constraints $y_{hri}(t)$.

Hormonal coupling. - Normin dominates. - Hazardin increases under collision risk. - Taskin is limited in human presence.

Cost. Low to medium.

5.4. Utility Functions and Costs: Robotic Instantiation

5.4.1. Activation Cost

Activation cost is decomposed as: $c_i = c_i^{cpu} + c_i^{energy} + c_i^{risk}$

Energexin modulates both the global budget $B(t)$ and local constraints $\chi_i(t)$. A simple instantiation is: $B(t) = B_0 (1 - h^{-Energexin}(t))$

5.4.2. Hormonal Utility

A generic form is: $u_i(h(t), s(t)) = \sum_k w_{i,k} \bar{h}_k(t) + \phi_i(s(t)) - \psi_i(\chi(t))$

where ϕ_i captures contextual relevance and ψ_i penalizes physical constraints.

5.5. Explicit Coupling with the Graph-Based Hormonal Engine

Agents operate on a distributed robotic platform with nodes i . Local hormones $h_{k,i}(t)$ evolve according to the reaction–diffusion dynamics of Section 4.4 [27].

This enables: - local reactivity, - controlled propagation, - stable arbitration, - direct incorporation of constraints $\chi_i(t)$.

Critical regimes may trigger local overrides, e.g.: if $h_{Hazardin,i}(t) \geq \tau_H$ then a_{avoid} is forced locally.

5.6. Anti-Oscillation Mechanisms: Hysteresis and Persistence

Beyond hysteresis, two additional mechanisms are enforced:

1. **Minimum activation duration (TTL):** once activated, a_i remains active for at least ΔT_{imin} .
2. **Cooldown:** after deactivation, a_i cannot be reactivated for ΔT_{icool} , except under high Hazardin.

These mechanisms are essential under noisy hormonal dynamics and Euler–Maruyama discretization.

5.7. Operational Summary of the Agent Library

Agent	Primary role	Dominant hormone	Cost c_i	Main output
Task Manager Agent	Mission planning	Taskin	high	goals, priorities
Navigation Agent	Nominal navigation	Taskin/ Normin	medium	navigation commands
Obstacle Avoidance Agent	Collision avoidance	Hazardin	medium	braking/ avoidance
Compliance Agent	Rule enforcement	Normin	low	constraints
Energy/Thermal Agent	Energy/thermal mgmt	Energexin	low	limitations
Recovery Agent	Stabilization/ recovery	Resiliencin	high	Recovery plan
Perception Health Agent	Sensor health	Hazardin/ Resiliencin	medium	reliability scores
Manipulation Agent	Manipulation	Taskin/ Normin	high	Arm commands
Human Interaction Agent	HRI regulation	Normin	medium	social constraints

5.8. Section Conclusion

The proposed typology is not a mere functional enumeration. It is **mathematically aligned** with: - primal–dual orchestration under budget $B(t)$, - stable hysteresis-based activation, - the graph-based hormonal engine with reaction–diffusion, cross-inhibition, local constraints $\chi_i(t)$, and bounded noise, - and symbolic engram encoding in memory [32].

The next section formalizes **robotic hormonal profiles and triggers**, including parameter calibration for $\alpha_{j,k}$, $\beta_{j,k}$, λ_k , $\gamma_{k\ell}$, D_k , and the construction of $\chi_i(t)$ within the SAI-UT+ framework.

6. HORMONAL PROFILES AND ROBOTIC TRIGGERS

This section formalizes the robotic hormonal profiles, triggering probes, and emission–propagation mechanisms implemented in S-AI-ROBOTICS [27]. It instantiates explicitly the equations introduced in Section 4, namely: (i) the saturated logistic emission $e_{j,k,i}(t)$, (ii) the graph-based reaction–diffusion dynamics with cross-inhibition and local constraints $\chi^i(t)$, and (iii) the Euler–Maruyama discretization with projection $\Pi_{[0,1]}$ [32]. The five canonical robotic hormones are: - **Taskin**: task pressure and execution urgency, - **Hazardin**: hazard and danger (collision, unexpected contact), - **Energexin**: energy, thermal, and power constraints, - **Resiliencin**: stabilization, recovery, and robustness, - **Normin**: mission and safety compliance (including HRI). We denote $k \in \mathcal{K} = \{T, H, E, R, N\}$ for Taskin, Hazardin, Energexin, Resiliencin, and Normin. Hormonal states evolve locally on a robotic graph $G=(V,E)$, whose nodes $i \in V$ represent physical or computational subsystems (sensors, compute units, mobile base, manipulators, local controllers). The normalized Laplacian L governs diffusion (Section 4.4).

6.1. Robotic Graph, Observables, and Normalization

6.1.1. Definition of the Robotic Graph $G=(V,E)$

The choice of graph topology follows the hardware and software architecture. Typical nodes include:

- Sensor nodes: lidar, camera, IMU,
- Actuator nodes: mobile base, manipulator,
- Compute nodes: CPU, GPU, microcontrollers,
- Logical nodes: state estimators, planners, safety monitors.

Edges $(i,j) \in E$ model physical couplings (base–arm), informational links (sensor–estimator), or functional dependencies (planning–control).

6.1.2. Local Robotic Observable Vector

Each node i provides a set of observables $z_{j,i}(t)$ processed by robotic glands g_j . We define the local observable vector: $\mathbf{o}_i(t) = (e_{\text{trk},i}(t), p_{\text{obs},i}(t), P_{\text{col},i}(t), \Delta\tau_i(t), b_{\text{att},i}(t), T_i(t), \ell_{\text{cpu},i}(t), v_i(t))$

where: - e_{trk} : tracking or task error, - p_{obs} : obstacle proximity, - P_{col} : estimated collision probability, - $\Delta\tau$: force–torque spikes or abnormal effort, - b_{att} : normalized battery level, - T : normalized temperature, - ℓ_{cpu} : compute load, - v : rule or compliance violation indicator.

Each observable is normalized to $[0,1]$ via a monotonic mapping $n(\cdot): \tilde{o}_{q,i}(t) = n(q, o_{q,i}(t)) \in [0,1]$ ensuring bounded and comparable hormonal emissions across modalities.

6.2. Robotic Glands and Emission Functions

6.2.1. Gland Typology

Robotic glands are organized into five main families:

- Task glands: errors, delays, backlog, progress,
- Hazard glands: obstacles, collisions, contact, instability,
- Energy/thermal glands: battery, power, temperature,
- Resilience glands: sensor degradation, instability, failure,

- Compliance glands: rule violations, restricted zones, HRI.

Each gland emits toward one or more hormones.

6.2.2. Saturated Logistic Emission

Following Section 4.3, for gland gj , hormone hk , and node i , the local emission is: $e_{j,k,i}(t) = \sigma(\alpha_{j,k}z_{j,i}(t) - \beta_{j,k})(1 - h_{k,i}(t))$

where: - $\alpha_{j,k}$ is a sensitivity parameter, - $\beta_{j,k}$ is an emission threshold, - $\sigma(x) = \frac{1}{1 + \exp(-x)}$, the factor $(1 - h_{k,i}(t))$ enforces saturation (diminishing returns).

This saturation is critical in robotics to prevent runaway Hazardin under sensor noise and to avoid Taskin overwhelming Normin under persistent task pressure [28].

6.2.3. Hormone-Level Emission Aggregation

For each hormone k and node i , the aggregated emission is:

$$e_{k,i}(t) = \sum_{j \in \mathcal{J}_k} w_{j,k} e_{j,k,i}(t)$$

where \mathcal{J}_k is the set of glands contributing to hormone k and $w_{j,k} \geq 0$ are fusion weights.

6.3. Graph-Based Reaction–Diffusion Hormonal Dynamics

6.3.1. Continuous-Time Model

Local hormones $h_{k,i}(t) \in [0, 1]$ evolve according to:

$$\dot{h}_{k,i}(t) = e_{k,i}(t) - \lambda_k h_{k,i}(t) - \sum_{\ell \neq k} \gamma_{k\ell} h_{\ell,i}(t) + D_k \sum_j L_{ij} h_{k,j}(t) - \chi_i(t) + \xi_{k,i}(t)$$

where: - λ_k is the decay rate, - $\gamma_{k\ell}$ encodes cross-inhibition, - D_k is the diffusion coefficient, - L is the normalized Laplacian, - $\chi_i(t)$ represents local physical and computational constraints, - $\xi_{k,i}(t)$ is bounded noise [27].

6.3.2. Cross-Inhibition Matrix Γ

The inhibition matrix $\Gamma = [\gamma_{k\ell}]$ satisfies: $\gamma_{k\ell} \geq 0$, $\gamma_{kk} = 0$

Canonical inhibitory relationships include: - Taskin inhibited by Hazardin and Normin, - Hazardin inhibits Taskin (safety priority), - Energexin inhibits Taskin (sobriety under constraints), - Normin inhibits Taskin and parts of Resiliencin, - Resiliencin inhibits Taskin during recovery and modulates Hazardin.

6.3.3. Diffusion and Robotic Rationale

The diffusion term $D_k \sum_j L_{ij} h_{k,j}$ enables controlled propagation, e.g.:

Hazardin rises at a sensor node \rightarrow propagates to control nodes,

Energexin rises at a thermal node → propagates to compute units,
Normin rises at an HRI node → propagates to navigation and manipulation.

Typical diffusion ordering: $DH > DN \gtrsim DE > DR \gtrsim DT$

6.4. Local Constraint Term $\chi_i(t)$

The local constraint term explicitly encodes physical and computational limitations:

$$\chi_i(t) = \omega_b \chi_i^{\text{batt}}(t) + \omega_T \chi_i^{\text{therm}}(t) + \omega_c \chi_i^{\text{cpu}}(t) + \omega_a \chi_i^{\text{act}}(t)$$

With

$$\chi_i^{\text{batt}}(t) = 1 - b_{\text{att},i}(t)$$

$$\chi_i^{\text{therm}}(t) = \sigma(\kappa_T(T_i(t) - T_{\text{max}}))$$

$$\chi_i^{\text{cpu}}(t) = \ell_{\text{cpu},i}(t)$$

$$\chi_i^{\text{act}}(t) = \sigma(\kappa_a(\| \tau_i(t) \| - \tau_{\text{max}}))$$

This term acts upstream of orchestration: high local constraints damp hormonal levels locally, indirectly discouraging the activation of costly agents via $ui(\cdot)$, $B(t)$, and activation thresholds [40].

6.5. Discretization and Noise: Euler–Maruyama with Projection

For SAI-UT+ simulations, the dynamics are discretized as: $h_{k,i}(t+1) = \Pi_{[0,1]}(h_{k,i}(t) + \Delta t \dot{h}_{k,i}(t) + \sigma_k \sqrt{\Delta t} \varepsilon_{k,i}(t))$

where $\varepsilon_{k,i}(t)$ is a zero-mean noise variable and $\Pi[0,1]$ projects onto $[0,1]$ [6].

Noise accounts for: - sensor latency, - measurement fluctuations, - environmental variability, - internal uncertainty.

6.6. Robotic Triggers (Probes) by Hormone

This subsection defines dominant triggers and observables per hormone.

6.6.1. Taskin: Task Pressure and Urgency

Typical triggers: - tracking error e_{trk} , - deadline delay $d(t)$, - backlog $b(t)$, - trajectory deviation.

$$z_{\text{task},i}(t) = \rho_1 e_{\text{trk},i}(t) + \rho_2 d_i(t) + \rho_3 b_i(t)$$

6.6.2. Hazardin: Danger and Risk

Triggers: - obstacle proximity p_{obs} , - collision probability P_{col} , - force–torque spikes $\Delta\tau$.

$$z_{\text{haz},i}(t) = \rho_4 p_{\text{obs},i}(t) + \rho_5 P_{\text{col},i}(t) + \rho_6 \Delta \tau_i(t)$$

6.6.3. Energexin: Energy and Thermal Constraints

Triggers: - low battery, - overheating, - high power draw, - compute overload.

$$z_{\text{ener},i}(t) = \rho_7 (1 - b_{\text{att},i}(t)) + \rho_8 T_i(t) + \rho_9 \ell_{\text{cpu},i}(t)$$

6.6.4. Resiliencin: Stabilization and Recovery

Triggers: - critical sensor degradation, - estimator divergence, - task failure, - detected oscillations.

$$z_{\text{res},i}(t) = \rho_{10} \text{fail}_i(t) + \rho_{11} \text{drift}_i(t) + \rho_{12} \text{unstable}_i(t)$$

6.6.5. Normin: Compliance and Safety

Triggers: - rule violation $v_i(t)$, - restricted zone entry, - human presence (HRI), - excessive speed.

$$z_{\text{norm},i}(t) = \rho_{13} v_i(t) + \rho_{14} \text{hri}_i(t) + \rho_{15} \text{speed}_i(t)$$

6.7. Hormonal Parameter Calibration

Key parameters include: - Emission: $\alpha_{j,k}, \beta_{j,k}, w_{j,k}$ - Decay: $\lambda_{k,t}$, - Inhibition: $\gamma_{k,t}$, - Diffusion: Dk , - Noise: σ_k , - Step size: Δ_t .

6.7.1. Constraint-Based Calibration

Calibration follows safe-by-construction constraints:

3. Emission boundedness: $0 \leq e_{j,k,i}(t) \leq 1$,
4. Decay: $\lambda_k \in (0, 1)$ in discrete time,
5. Moderate inhibition to avoid global extinction,
6. Diffusion tuned to graph degree
7. Noise lower for compliance hormones, higher for Taskin in exploration [39].

6.7.2. Memory-Guided Calibration

Contextual memory (Section 4.7) stores engrams $\langle h(t), \mathbf{x}(t), a(t), o(t) \rangle$, enabling reuse of parameter sets associated with validated regimes [32].

6.8. Coupling Hormones to Parsimonious Orchestration

Hormones influence orchestration through:

8. Agent utility functions $u_i(h(t), s(t))$,
9. The global budget $B(t)$, particularly via Energexin.

A consistent instantiation is:

$$u_i(h(t), s(t)) = \sum w_{i,k} \bar{h}_k(t) + \phi_i(s(t)) - \psi_i(\chi(t))$$

$$B(t) = B_0(1 - \bar{h}_{\text{Energexin}}(t))$$

This coupling ensures that under constrained regimes (high Energexin), costly agents are suppressed in favor of safe and frugal behaviors [40].

6.9. Section Summary

This section has defined: - the robotic graph G and normalized observables, - robotic glands and saturated logistic emissions, - reaction–diffusion dynamics with inhibition, diffusion, constraints $\chi_i(t)$, and noise, - Euler–Maruyama discretization with projection, - hormone-specific triggers and calibration principles. The next section formalizes parsimonious orchestration and stability guarantees, building on primal–dual dynamics, hysteresis, and memory-assisted regulation.

7. PARSIMONIOUS ORCHESTRATION AND STABILITY ANALYSIS

This section formalizes the parsimonious orchestration mechanism of S-AI-ROBOTICS and provides a rigorous analysis of its stability properties. The orchestration layer bridges the hormonal dynamics defined in Section 4 with the execution of specialized robotic agents described in Section 5 [32]. Its role is to ensure sparse activation, robustness under uncertainty, and stable behavior selection in the presence of physical, energetic, and safety constraints [6].

7.1. Orchestration as Constrained Optimization under Hormonal Regulation

At each decision cycle t , S-AI-ROBOTICS selects a subset of specialized agents through a continuous activation vector $\mathbf{x}(t) \in [0,1]^N$. The orchestration problem is defined as:

$$\max_{\mathbf{x}(t)} \sum_{i=1}^N x_i(t) u_i(h(t), s(t)) \quad \text{subject to} \quad \sum_{i=1}^N x_i(t) c_i \leq B(t)$$

where: - $u_i(h(t), s(t))$ is the hormone-modulated utility of agent ai , - c_i is the activation cost, - $B(t)$ is the available global budget, modulated by Energexin. This formulation enforces context-aware sparsity: only agents with sufficiently high utility relative to their cost and the available budget are activated [4].

7.2. Primal–Dual Dynamics for Online Orchestration

To solve the constrained optimization online, S-AI-ROBOTICS employs a primal–dual dynamic system:

$$\begin{aligned} \dot{x}_i(t) &= \eta_i \Pi_{[0,1]}(u_i(h(t), s(t)) - \lambda(t) c_i - \rho_a x_i(t)) \\ \dot{\lambda}(t) &= \eta_\lambda (\sum_{i=1}^N x_i(t) c_i - B(t))_+ \end{aligned}$$

where: - $\lambda(t)$ is the Lagrange multiplier associated with the budget constraint, - $\rho_a > 0$ is a regularization coefficient enforcing sparsity, - η_i and η_λ are step-size parameters, - $\Pi_{[0,1]}$ projects onto the admissible activation interval. This dynamic ensures that: 1. The budget constraint is satisfied asymptotically. 2. Agents with low marginal utility are progressively deactivated. 3. The orchestration adapts smoothly to changes in hormonal states [6].

7.3. Coupling with Hormonal Reaction–Diffusion Dynamics

The orchestration variables $\mathbf{x}(t)$ are indirectly coupled to the hormonal dynamics via the utility functions $u_i(h(t),s(t))$. Hormonal states $h_{k,i}(t)$ evolve according to the reaction–diffusion system defined in Section 4.4, including cross-inhibition, diffusion, and local constraints $\chi_i(t)$ [27].

The coupling operates at two levels:

10. **Global coupling:** aggregated hormones $h^{-k}(t)$ modulate utilities u_i and the global budget $B(t)$.
11. **Local coupling:** critical local hormonal thresholds (e.g., $h_{\text{Hazardin},i}(t)$) can override orchestration decisions and force agent activation or deactivation.

This dual coupling ensures that safety-critical information propagates rapidly while preserving global coherence [38].

7.4. Hysteresis-Based Discretization and Execution Stability

The continuous activation values $x_i(t)$ are mapped to discrete execution decisions using hysteresis thresholds:

$$a_i(t^+) = \begin{cases} 1 & \text{if } x_i(t) \geq \theta_i^\uparrow \\ 0 & \text{if } x_i(t) \leq \theta_i^\downarrow \\ a_i(t^-) & \text{otherwise} \end{cases}$$

With $\theta_i^\uparrow > \theta_i^\downarrow$.

This mechanism prevents rapid ON/OFF switching when $x_i(t)$ fluctuates near activation thresholds due to noise in hormonal dynamics or sensory inputs [7].

7.5. Time-Scale Separation and Singular Perturbation Argument

A key property of S-AI-ROBOTICS is the separation of time scales between:

- fast perceptual and glandular emissions,
- intermediate hormonal reaction–diffusion dynamics,
- slower orchestration and agent execution cycles.

Let $\tau_s \ll \tau_h \ll \tau_o$ denote the characteristic time scales of sensing, hormonal evolution, and orchestration, respectively. Under this separation, the hormonal system can be approximated as quasi-stationary during a single orchestration update [6].

This separation justifies the use of continuous primal–dual dynamics and supports stability analysis via singular perturbation arguments, ensuring that fast hormonal fluctuations do not destabilize agent selection.

7.6. Stability of the Primal–Dual Orchestration Loop

Consider the Lyapunov candidate function:

$$V(\mathbf{x}, \lambda) = \sum_{i=1}^N \frac{1}{2} (x_i - x_i^*)^2 + \frac{1}{2} (\lambda - \lambda^*)^2$$

where $(\mathbf{x}^*, \lambda^*)$ denotes an equilibrium point of the primal–dual system.

Under standard assumptions: - bounded utilities u_i , - positive costs c_i , - sufficiently small step sizes η_i, η_λ ,

the time derivative V is non-positive, implying convergence to a stable equilibrium set[6]. This guarantees that orchestration decisions remain bounded and converge under stationary hormonal regimes.

7.7. Robustness under Noisy Hormonal Dynamics

Hormonal dynamics are subject to bounded noise $\xi_{k,i}(t)$ and stochastic discretization (Euler–Maruyama). The orchestration loop remains robust due to:

12. Projection operators $\Pi[0,1]$ limiting activation values.
13. Hysteresis preventing chattering at execution level.
14. Regularization $\rho ax_i(t)$ damping oscillations.

As a result, orchestration behaves as a low-pass filter over hormonal fluctuations, preserving decision stability under uncertainty [39].

7.8. Safety Overrides and Hard Constraints

Certain hormonal regimes impose non-negotiable safety constraints. For instance:

$$h_{\text{Hazardin},i}(t) \geq \tau_H \quad \Rightarrow \quad a_{\text{avoid}}(t) = 1, a_{\text{task}}(t) = 0$$

Such overrides are modeled as hard constraints that supersede the optimization process. They ensure that safety and compliance are always prioritized over performance [38].

7.9. Computational Complexity and Scalability

Let N be the number of agents and K the number of hormones. Each orchestration update has complexity:

- $O(N)$ for utility evaluation and primal updates,
- $O(K|E|)$ for hormonal diffusion over the graph.

This linear scaling makes S-AI-ROBOTICS suitable for embedded and real-time robotic platforms [40].

8. GENERAL CONCLUSION

This paper introduced S-AI-ROBOTICS, a bio-inspired robotic intelligence architecture grounded in the principles of Sparse Artificial Intelligence and designed explicitly for embodied, safety-critical, and resource-constrained robotic systems. Departing from monolithic and continuously active control paradigms, the proposed framework models robotic cognition as a regulated and context-aware orchestration of specialized agents. At the core of the architecture lies a Robo-Meta Agent that enforces selective activation through constrained optimization and hysteresis-based dynamics. Rather than activating all subsystems indiscriminately, S-AI-ROBOTICS ensures that only a minimal subset of agents is engaged at any given time, preserving computational and energetic resources while maintaining coherent task execution. In this sense, parsimony is elevated from a heuristic design choice to a formal regulatory principle embedded in the decision process itself. A central contribution of this work is the introduction of an artificial hormonal signaling layer tailored to embodied robotics. Inspired by neuroendocrine systems, the hormonal model regulates behavioral priorities through bounded emission, diffusion, inhibition, and decay mechanisms. Hormones do not encode actions directly; instead, they modulate activation thresholds and resource budgets, enabling smooth arbitration between competing objectives such as urgency, safety, stability, and energy conservation. This regulatory layer provides a principled mechanism for conflict resolution and behavioral coherence under uncertainty. The architecture is further strengthened by a symbolic and contextual memory subsystem that stores behavioral engrams linking hormonal context, activated agents, actions, and outcomes. This structured memory supports rapid adaptation through recall of validated strategies and enhances interpretability by preserving causal traces of decision-making. Explainability thus emerges natively from the architecture rather than being imposed as a post-hoc analytical layer. Collectively, hormonal regulation, sparse orchestration, and symbolic memory form an integrated framework for embodied intelligence that is selective, stable, and inherently interpretable. S-AI-ROBOTICS demonstrates that adaptive robotic systems can be designed around regulatory principles that balance performance, safety, and energy efficiency without sacrificing transparency. Future work will extend this foundation toward large-scale real-world deployment, formal stability analysis under hardware constraints, and deeper integration with learning-based components operating under strict regulatory supervision. By embedding regulation, frugality, and explainability at the architectural level, S-AI-ROBOTICS provides a principled pathway toward next-generation robotic systems capable of operating safely and efficiently in complex, uncertain environments.

REFERENCES

- [1] H.-G. Han, S. Fu, H. Sun, and J. Qiao, "Data-driven model predictive control for nonlinear systems with stochastic sampling interval," *IEEE Transactions on Industrial Electronics*, vol. 69, no. 8, pp. 8213–8223, 2022.
- [2] I. Moreno-Caireta, J. B. A. Reyes, and R. Carelli, "Model predictive control for a mecanum-wheeled mobile robot under constraints," *IFAC-PapersOnLine*, vol. 54, no. 10, pp. 283–288, 2021.
- [3] J. Rakhmatillaev, "An integrative review of control strategies in robotics," *Robotics and Autonomous Systems*, vol. 173, Art. no. 104594, 2025.
- [4] E. F. Camacho and C. Bordons, *Model Predictive Control*, 2nd ed. London, U.K.: Springer, 2019.
- [5] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*. London, U.K.: Springer, 2010.
- [6] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton, NJ, USA: Princeton Univ. Press, 2008.
- [7] M. Wang, Y. Li, and Z. Chen, "Hybrid control systems in robotics: A survey," *IEEE Access*, vol. 11, pp. 45612–45630, 2023.
- [8] L. Grüne and J. Pannek, *Nonlinear Model Predictive Control*. Cham, Switzerland: Springer, 2017.
- [9] A. K. Shakya, S. M. Shrestha, and R. P. Bhattarai, "Reinforcement learning algorithms: A survey," *Expert Systems with Applications*, vol. 213, Art. no. 119038, 2023.

- [10] D. Han, Y. Chen, and H. Wang, "A survey on deep reinforcement learning algorithms for robotic manipulation tasks," *Sensors*, vol. 23, no. 7, Art. no. 3762, 2023.
- [11] S. Milani, L. Longo, and A. Alonso, "Explainable reinforcement learning: A survey and taxonomy," *ACM Computing Surveys*, vol. 56, no. 4, Art. no. 89, 2024.
- [12] Y. Lan, J. Liu, and X. Zhang, "Reinforcement learning for optimal decision-making in autonomous systems: A survey," *IET Cyber-Physical Systems*, vol. 7, no. 1, pp. 1–15, 2025.
- [13] J. Kober and J. Peters, "Reinforcement learning in robotics: A survey," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [14] L. Brunke, M. Greeff, A. W. Hall, and J. Peters, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Reviews in Control*, vol. 54, pp. 34–52, 2022.
- [15] A. García and M. Fernández, "A survey on safe reinforcement learning," *Journal of Machine Learning Research*, vol. 26, pp. 1–58, 2025.
- [16] T. Haarnoja et al., "Soft actor-critic algorithms and applications," *arXiv preprint arXiv:1812.05905*, 2018.
- [17] K. Arulkumaran et al., "Deep reinforcement learning: A brief survey," *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 26–38, 2017.
- [18] S. Levine et al., "End-to-end training of deep visuomotor policies," *Journal of Machine Learning Research*, vol. 17, pp. 1–40, 2016.
- [19] M. Iovino et al., "A survey of behavior trees in robotics and AI," *Robotics and Autonomous Systems*, vol. 153, Art. no. 104096, 2022.
- [20] S. Gugliermo et al., "Evaluating behavior trees for robotic task execution," *Robotics and Autonomous Systems*, vol. 176, Art. no. 104628, 2024.
- [21] A. Colledanchise and P. Ögren, *Behavior Trees in Robotics and AI*. Piscataway, NJ, USA: IEEE Press, 2018.
- [22] R. A. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal of Robotics and Automation*, vol. 2, no. 1, pp. 14–23, 1986.
- [23] L. Sentis and O. Khatib, "Synthesis of whole-body behaviors through hierarchical control," *The International Journal of Robotics Research*, vol. 24, no. 7, pp. 505–525, 2005.
- [24] M. Ghallab, D. Nau, and P. Traverso, *Automated Planning: Theory and Practice*. San Francisco, CA, USA: Morgan Kaufmann, 2004.
- [25] J. Bohg et al., "Task and motion planning for robotic manipulation," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 475–501, 2018.
- [26] H. Van Brussel et al., "Reference architecture for holonic manufacturing systems," *Computers in Industry*, vol. 37, no. 3, pp. 255–274, 1998.
- [27] J. Timmis et al., "An adaptive neuro-endocrine system for robotic systems," *Artificial Intelligence*, vol. 173, no. 9–10, pp. 983–1018, 2009.
- [28] M. Maroto-Gómez et al., "A motivational model based on artificial biological functions for autonomous robots," *Autonomous Robots*, vol. 47, pp. 1–21, 2023.
- [29] M. Maroto-Gómez et al., "Human-robot pair-bonding from a neuroendocrine robotics perspective," *Robotics and Autonomous Systems*, vol. 174, Art. no. 104561, 2024.
- [30] R. Pfeifer and J. Bongard, *How the Body Shapes the Way We Think*. Cambridge, MA, USA: MIT Press, 2007.
- [31] S. Doncieux et al., "Evolutionary robotics: What, why, and where to," *Artificial Intelligence*, vol. 173, no. 7–8, pp. 1257–1270, 2015.
- [32] S. Slaoui, "S-AI: A sparse artificial intelligence system orchestrated by a hormonal MetaAgent and context-aware specialized agents," *International Journal of Formal Methods & Research*, vol. 1, no. 2, pp. 1–18, 2025.
- [33] S. Slaoui, "Bio-inspired hormonal modulation and adaptive orchestration in S-AI-GPT," *International Journal of Artificial Intelligence & Applications*, vol. 16, no. 4, 2025.
- [34] S. Slaoui, "S-AI-NET: A sparse AI framework for adaptive and parsimonious autonomous networking," *Research Square*, 2025.
- [35] L. Longo, "Explainable artificial intelligence (XAI) 2.0: A manifesto," *Information Fusion*, vol. 98, pp. 101–121, 2024.
- [36] M. Anjomshoa et al., "Explainable artificial intelligence for robotics: A survey," *IEEE Access*, vol. 11, pp. 34421–34445, 2023.

- [37] M. Diehl and K. Ramirez-Amaro, "Why did I fail? A causal-based method to find explanations for robot failures," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8261–8268, 2022.
- [38] A. Ghalamzan et al., "Safe and explainable robotics: A survey," *Robotics and Autonomous Systems*, vol. 167, Art. no. 104442, 2023.
- [39] L. Brunke et al., "Safe learning for control," *Annual Reviews in Control*, vol. 54, pp. 1–20, 2022.
- [40] A. Gupta et al., "Energy-efficient robotics: A systematic literature review," *IEEE Access*, vol. 12, pp. 8890–8920, 2024.

AUTHORS

Said Slaoui is a professor at Mohammed V University in Rabat, Morocco. He graduated in Computer Science from University Pierre and Marie Curie, Paris VI (in collaboration with IBM France), 1986. He has over 40 years of experience in the fields of AI and Big Data, with research focused on modular architectures, symbolic reasoning, and computational frugality. His recent work introduces the Sparse Artificial Intelligence (S-AI) framework, which integrates bio-inspired signaling and agent-based orchestration. He has published numerous scientific papers in international journals and conferences, and actively contributes to the development of sustainable and explainable AI systems.

