A unified control-theoretic architecture for emotion, cognition and adaptive behaviour in biological and artificial agents

Rocco Van Schalkwyk[1], Denise Cook[1], Alireza Dehbozorgi[1], Carlos E. Alvarez[2]

[1]Xzistor LAB, Bristol, United Kingdom

[2]Departments of Pediatrics and Veterinary Clinical Sciences, The Ohio State University, Columbus, Ohio, USA

Correspondence: rocco.vanschalkwyk@xzistor.com

Abstract

How sensation, emotion and cognition arise from a unified neural architecture is a central problem in neuroscience and artificial intelligence. Here we present the Xzistor Mathematical Model of Mind (XMMM), a control-theoretic architecture in which all adaptive behaviour follows from a single principle: the brain minimises homeostatic and allostatic deprivation, experienced as body-located somatosensory emotions. Five interoperating algorithms transform multi-variable homeostatic complexity into valenced emotional representations, enabling the executive to act on felt experience rather than raw physiological data. Implemented in a virtual agent and a physical robot, the architecture produces reinforcement learning, goal-directed navigation, fear conditioning, social attachment and flexible reasoning without externally specified reward functions. A parametric lesion study establishes a dissociation between reactive associative behaviour and directed inductive inference. Post-hoc functional convergence with interoceptive, episodic memory and action-selection systems constitutes independent biological validation. Five quantitative predictions are specified with explicit falsification conditions. The XMMM delivers computable accounts of emotion and cognition—and a neuro-symbolic pathway toward artificial general intelligence.

Introduction

Understanding how physiological need states generate felt emotional experience, motivate behaviour and shape learned associations is a central problem across neuroscience, cognitive science and artificial intelligence. In biological organisms, the regulation of internal variables—energy balance, hydration, temperature, social contact—is not merely metabolic background: it is the primary organising principle of adaptive behaviour. Contemporary frameworks converge on the view that the brain operates as a predictive control system whose function is to maintain internal homeostasis while interacting with the external environment, and that felt emotional states—not abstract reward signals—constitute the primary motivational currency of biological cognition.[1–3] Despite this convergence, a unified mechanistic framework that translates these principles into a

formally specified, implemented and empirically tested computational architecture has remained absent.

Existing cognitive architectures—ACT-R,[4] SOAR[5] and Global Workspace Theory[6]—have advanced understanding of memory retrieval, problem-solving and attentional broadcasting, but treat emotion and cognition as separable subsystems and have not been demonstrated in physical agents under naturalistic drive-motivated conditions. Computational neuroscience models of hypothalamic survival circuits[7] and action-outcome learning[8] have clarified how bodily states influence decision-making, but do not provide implementable formal specifications testable in autonomous agents. Reinforcement learning (RL) generates goal-directed behaviour from reward signals, but those signals are typically defined externally by task designers rather than generated intrinsically through homeostatic regulation—a fundamental difference from the biological case.[9–11] The Free Energy Principle frames cognition as prediction-error minimisation and subsumes homeostatic regulation,[2,12] but its predictions are formulated at the level of computational principles rather than measurable physiological units. No existing framework simultaneously implements emotion as a computable body-located somatosensory quantity, derives it from homeostatic drive dynamics, and demonstrates the emergence of goal-directed behaviour, fear conditioning, social bonding and flexible reasoning from a single minimal architecture.

A particular gap concerns the relationship between homeostatic emotion and higher cognition. The capacity to generalise learned behaviour to novel contexts and to handle abstract concepts such as planning has generally been attributed to specialised modules distinct from basic motivated behaviour. We propose that these capacities result naturally from drive-motivated associative retrieval without additional substrate. Control theory offers a principled account of subjective emotional experience: the transformation of multi-variable homeostatic complexity into single body-located emotional representations is the necessary informational architecture for a real-time executive controller operating under biological time constraints.

Here we provide the first comprehensive account of the XMMM. Derived from control-theoretic first principles rather than neuroanatomy, the XMMM reproduces observable autonomous agent behaviours from a minimal initial endowment. We show that reinforcement learning, goal-directed navigation, fear conditioning, social bonding and parametric deliberative flexibility emerge without task-specific reward engineering, pre-trained components or explicitly programmed higher-level behaviours. Post-hoc functional convergence with interoceptive, episodic memory and action-selection systems constitutes independent corroboration of functional completeness, and the model yields specific, quantifiable predictions testable in human and rodent cohorts.

Results

Model architecture

The XMMM models the brain as a closed-loop multivariable adaptive control system comprising five functional algorithms—Sensing (S), Drive (D), Reflex (R), Association (A) and Motion (M)—coordinated by a Linking Algorithm acting as the executive controller (Fig. 1). Each drive i monitors at least one physiological control variable $\mathrm{CV}_i(t)$ against a homeostatic set-

point SP_i, generating a normalised error signal (equation 1). This signal is transformed into a body-located Deprivation Emotion (DE; equation 2) when rising, or a Satiation Emotion (SE; equation 3) encoding the rate of recovery when falling. The executive controller receives only these body-located representations—never the underlying physiological values—and selects the Prime Drive (PD) as the arg max of all signed drive values (equation 4). All learning, memory retrieval and action selection are governed by these felt emotional signals. Full specification, including equations, pseudocode and complexity analysis, is provided in the Methods section and Supplementary Document S0 (Methods Extended).

$$\mathrm{ES}_i(t) = \frac{|\mathrm{CV}_i(t) - \mathrm{SP}_i|}{\mathrm{CV}_{\max,i}} \in [0,1] \tag{1}$$

$$\mathrm{DE}_i(t) = -\mathrm{ES}_i(t) \in [-1,0] \quad \text{(aversive; felt as bodily need)} \tag{2}$$

$$\mathrm{SE}_i(t) = \left|\frac{d\mathrm{ES}_i}{dt}\right| \in [0,1] \quad \text{(appetitive; felt as bodily relief)} \tag{3}$$

$$\mathrm{PD}(t) = \arg\max_i |D_i(t)| \tag{4}$$

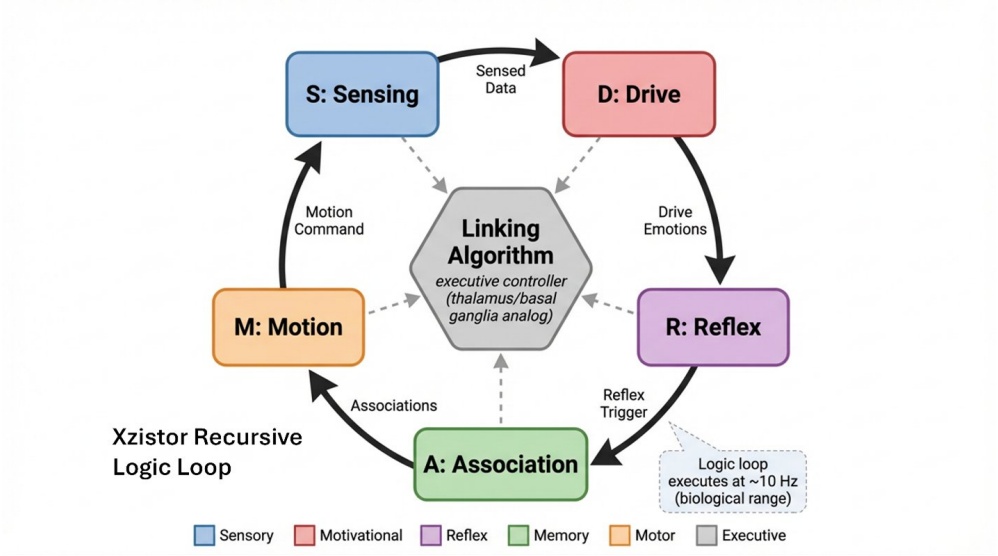

Figure 1 The XMMM recursive logic loop architecture. Five functional algorithms (Sensing, Drive, Reflex, Association, Motion) coordinated by the Linking Algorithm acting as the executive controller. Arrows show information flow between algorithms; dashed arrows indicate Linking Algorithm oversight. The loop executes at approximately 10 Hz in basic implementations, with frequency decreasing as the Association Database grows, consistent with the $\mathcal{O}(|\mathcal{A}| \log |\mathcal{A}|)$ retrieval complexity prediction.

Experimental context

The XMMM was implemented and tested under naturalistic conditions— variable terrain, lighting, sounds, starting orientations and hardware—because variable conditions are the appropriate

test of an architecture claiming to model biological cognition. Because each trial begins from a distinct agent state, trials are not statistically equivalent and results are reported as learning trajectories and internal state metrics, which the architecture's full introspectability uniquely affords. Full experimental protocols and results are documented in Supplementary Note S1.

Reinforcement learning in a virtual agent

The architecture was first implemented in a virtual agent ('Simmy') operating in a simulated environment with objects inducing positive or negative drive effects. The agent possessed eleven active drives— hunger, thirst, fatigue, temperature regulation, nausea, fear and anger—and started with no prior knowledge or pre-programmed behaviours. As the agent explored, interactions with reward objects generated satiation emotions that reinforced the corresponding experience associations, incrementing their Impact Factor (IF)—a scalar encoding emotional salience, recency and recall frequency. Associations consistently preceding satiation accumulated progressively larger IF values and became increasingly likely to be retrieved in future action selection.

The Reward-based Backpropagation (RBP) reinforcement signal is intrinsic and homeostatic: $\mathrm{SE}_i(t)$ is generated internally from the rate of change of the drive error signal without any externally specified reward function. Pre-coded virtual-tutor subroutines temporarily guide motor output when specific drive thresholds are met; because these trajectories pass through the same RBP channel as autonomous actions, the transition from instinct to skill is architecturally seamless (Fig. 2). As associative memory accumulated, behaviour shifted from exploratory wandering toward goal-directed navigation, with time-to-reward decreasing across training sessions and paths becoming progressively more direct.

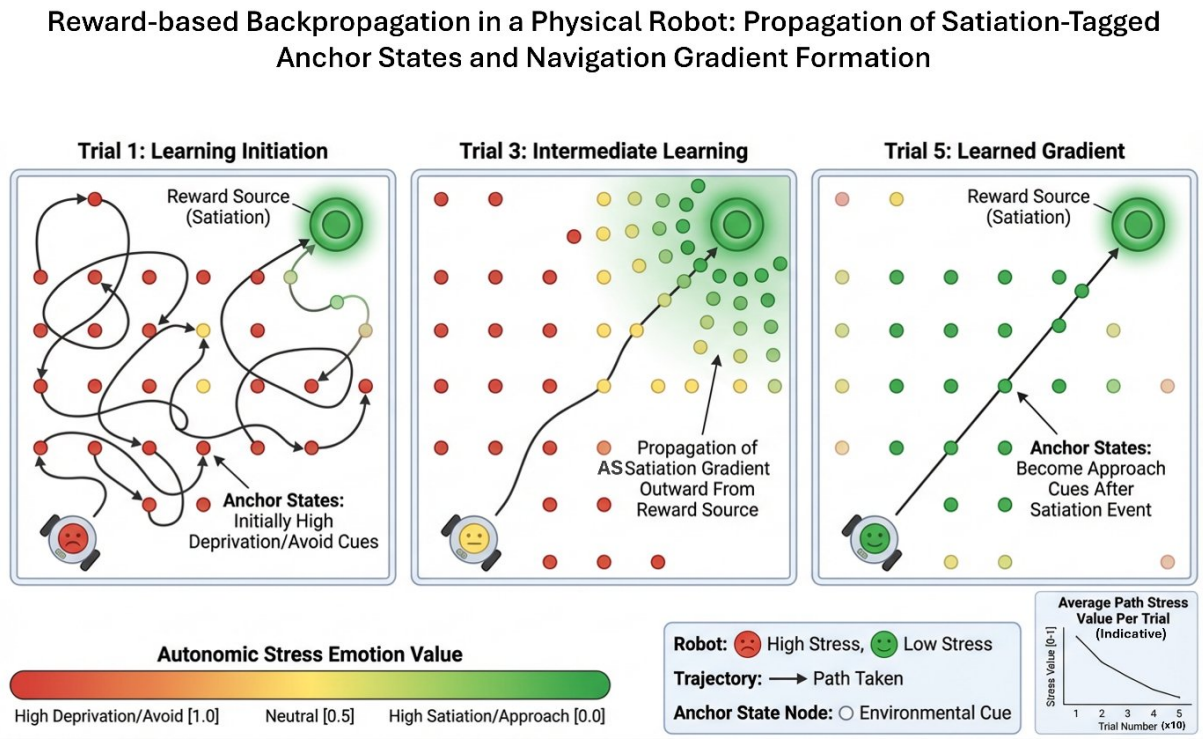

Figure 2 Reward-based backpropagation (RBP) in a physical robot. Sequential snapshots showing the propagation of satiation-tagged Anchor States—stored compound sensory and drive-emotional context markers—from the satiation source outward through the learning environment, forming a navigation gradient. Anchor States (dots) become tagged with Autonomic Stress (AS) satiation (relief) coupled to the Prime Drive (PD). The inset chart (indicative) shows average path stress declining across trials as the navigation gradient matures.

Goal-directed navigation in a mobile robot

To evaluate generalisation to embodied agents, the XMMM was implemented in a mobile robot ('Troopy') with twelve active drives. Crucially, the hunger drive was implemented as a time-based proxy homeostat—a numerical control variable $\mathrm{CV}_{\mathrm{hunger}}(t)$ rising independently of any physical sensor—illustrating the Epistemic Isolation Principle: the executive acts on what it feels, not on the nature of the underlying control variable. The same executive logic therefore governs physiological and purely numerical homeostats alike, demonstrating substrate independence in practice.

Over successive trials the robot learned to navigate toward the reward location with decreasing path length and tutor-assistance time. Several behaviours emerged from the architecture without bespoke programming: conditioned fear responses to pain-associated objects,[13] stress-proportional facial expressions, daydreaming, and reward-source route imagination. Together these demonstrate that a compact set of control-theoretic algorithms generates a qualitatively rich behavioural repertoire.

Problem-solving through associative threading

Beyond reinforcement learning, the architecture supports goal-directed problem-solving through Directed Threading—a drive-motivated systematic traversal of the Association Database in

which candidate associations are ranked by drive-contextualised Impact Factor and evaluated serially. Troopy was trained to obtain food by pressing a specific coloured button on a four-button panel; the reward-producing button was then changed without further training or environmental cues.

When the learned action failed, rising hunger crossed the activation threshold ($\theta_{\text{act}} = 0.10$), triggering Directed Threading. The agent retrieved stored associations from related contexts and applied their motor commands to the modified panel, testing candidate solutions drawn from prior experience. Within a small number of search cycles, Troopy identified the new reward-producing button through inductive inference and achieved satiation.

This result distinguishes the XMMM sharply from standard RL: the driving force is homeostatic deprivation rather than an extrinsic reward signal, and the inference mechanism is drive-motivated associative search rather than policy gradients. The same substrate scales to abstract domains because the architecture is agnostic to the origin of the unsatisfied association (Supplementary Document S3, Sections 1.1–1.2).

Convergent validation of drive architecture

Two biological constraints, identified through observation of physiologically implausible agent behaviour, were independently incorporated into the Drive Algorithm and subsequently found to match established biology. First, a gustatory delay modelling the known oropharyngeal satiation latency in mammals was required to eliminate high-frequency approach-retreat oscillations. Second, fractional inhibitory coupling from the thirst error signal onto hunger (coefficient 0.10) reflected known hypothalamic cross-inhibition. Both constraints were derived from observed agent behaviour and converge independently on established biological mechanisms, providing empirical validation of the Drive Algorithm's control-theoretic structure (Supplementary Document S2, Section 1.1).

Social bonding and attachment-figure differentiation

To test whether the architecture supports primary social emotion independent of secondary reinforcement, a Social Bonding (SB) drive was implemented in Troopy as a dedicated homeostatic loop. The bonding stimulus was a high-contrast blue colour panel on a Mobile Feeder frame; the feeder provided no food, water, or any other drive reinforcement, directly excluding the cupboard-love mechanism. Physical proximity was registered by a downward-facing sensor detecting a green floor panel beneath the feeder (Fig. 3**f**). The drive error signal ES_{SB} accumulates linearly during absence of the bonding stimulus and decays exponentially upon recognition—detected by Correlation Sensitivity $\text{CS}_{\text{SB}} \geq \theta_{\text{CS}} = 0.90$—combined with floor contact:

$$\text{ES}_{\text{SB}}(t) = \min\!\big(1,\ \kappa_{\text{dep}} \cdot t_{\text{abs}}(t)\big), \quad \text{CS}_{\text{SB}}(t) < \theta_{\text{CS}} \tag{5}$$

$$\frac{d\text{ES}_{\text{SB}}}{dt} = -\kappa_{\text{sat}} \cdot \text{ES}_{\text{SB}}(t), \quad \text{CS}_{\text{SB}}(t) \geq \theta_{\text{CS}} \text{ and floor contact} \tag{6}$$

Four classes of emergent, unprogrammed behaviour were observed. In the high-dominance configuration, Troopy immediately maintained contact with the feeder, reproducing neonatal proximity maintenance (Fig. 3**f**); when a pain-triggering zone was placed beyond the contact panel, the dominant SB drive maintained approach behaviour despite repeated aversive stimulation (Fig. 3**h**). In the naturalistic configuration, departure distances increased progressively as the agent exploited the amplified satiation magnitude of delayed reunion (Fig. 3**g**)—a computational realisation of Bowlby's secure-base concept.[10] Configuring the feeder also as a food source compounded the Impact Factor of bonding-associated states, mirroring the enhanced attachment to caregivers who satisfy multiple concurrent needs.

The biological basis for modelling social bonding as a homeostatic drive is supported by hypothalamic medial preoptic neurons that encode social need and social satiety in direct analogy to hunger-governing AgRP/POMC populations,[29,30] with amygdala stress coupling generating the aversive phenomenology of loneliness and an oxytocin–dopamine pathway implementing satiation.[31] The full eight-component circuit correspondence is in Supplementary Document S2, Section 1.4.

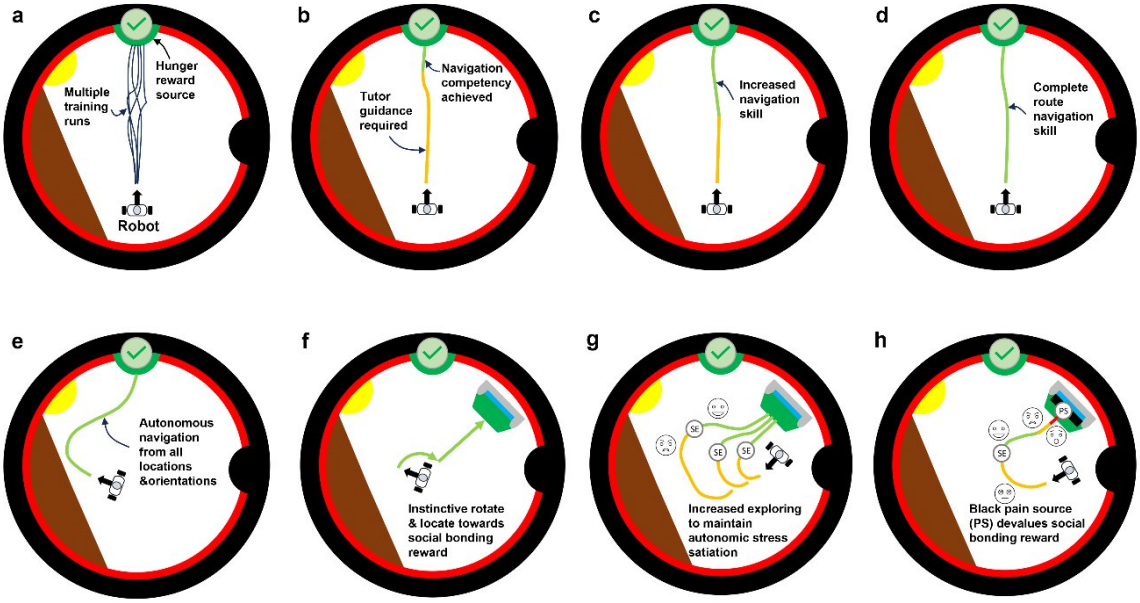

Figure 3 Navigation skill and social bonding behaviours. Eight plan-view snapshots of Troopy in a circular learning arena. **a**–**e**, Sequential learning of a navigation competency toward a hunger reward source through tutor-guided training runs. **f**, Instinctive approach toward a social bonding reward source (blue colour panel; green floor panel registers comfort contact). **g**, Increased exploration to sustain Autonomic Stress (AS) satiation events (SE), producing safe-base exploration. **h**, Devaluation of the bonding source after introduction of a pain source (PS, black band). Full protocol in Supplementary Note S1, Section 5.4.

A parametric dissociation of reactive and deliberative cognition

To isolate the contribution of Directed Threading from associative learning, threading was programmatically disabled in Troopy while all other components—Drive, Reflex and Association Algorithms and RBP—remained fully active. With threading disabled, the robot continued to

learn through RBP and, when presented with a high-fidelity match to a trained context, immediately retrieved and executed the correct motor sequence. The deficit appeared exclusively in partially familiar, untrained configurations: the threading-disabled robot ceased goal-directed navigation, became progressively more deprived, and defaulted to a pre-programmed hunger distress reflex, mirroring the cue-dependent, reactive profile of non-human animals in novel environments. This result establishes a clean dissociation between associative recognition-execution (intact without threading) and directed inductive inference (absent without threading). Full experimental protocol and biological correspondence analysis are in Supplementary Document S1 (Section 5.5) and Supplementary Document S2 (Section 2.6) respectively.

Discussion

These results show that reinforcement learning, goal-directed navigation, fear conditioning, problem-solving and social preference can all emerge from a single control-theoretic principle—minimise homeostatic and allostatic deprivation experienced as body-located somatosensory emotion—operating through a five-algorithm recursive logic loop, without externally specified reward functions, pre-trained components or explicitly programmed higher-level behaviours. The emergent fear responses, facial expressions and daydreaming arose from mechanisms already required for basic homeostatic regulation, not additional programming. The parametric lesion study further shows that the shift from reactive to deliberative cognition is a single-parameter variation within a conserved architecture, consistent with comparative neuroanatomical evidence for a prefrontal elaboration gradient across mammalian species.

The architecture's central claim concerns why subjective emotional experience exists at all. A drive such as thirst integrates plasma osmolality, blood volume and angiotensin II—none of which the executive can directly interpret. The Body Map performs sub-executive data compression: the Drive Algorithm absorbs all physiological non-linearity and delivers a single valenced scalar—a felt bodily state—to the executive. The executive does not process plasma osmolality; it experiences thirst. This transformation is the necessary informational interface between subliminal physiological complexity and a real-time controller, constituting the XMMM's formal account of why subjective emotional experience is an architectural requirement rather than an epiphenomenon. This account is consistent with the insular cortex's posterior-to-anterior interoceptive hierarchy[1,13] and implements somatic marker theory[23] at a mechanistic, computable level.

Generalisation emerges from homeostatic urgency: when Directed Threading fails to find an exact Anchor State match, drive pressure reduces the Correlation Sensitivity (CS) threshold and the architecture applies the highest-IF partial match on a trial-and-error basis, reinforcing successful transfers through RBP. Abstraction follows from the same mechanism: procedural sequences that reliably reduce time-to-satiation across diverse contexts accumulate high IF values and become retrievable via inductive threading in novel domains. A strategic plan and a recipe are structurally identical in the XMMM: both are high-IF Anchor States encoding efficient pathways to homeostatic relief (Supplementary Document S3, Sections 1.1–1.2).

Although derived from control-theoretic principles rather than neuroanatomy, the XMMM's

functional components show post-hoc convergence with known biological structures: the Body Map with the insular cortex's interoceptive hierarchy,[1,13] the Drive Algorithm with hypothalamic homeostatic regulation,[14] the Association Algorithm with hippocampal memory indexing,[15,16] the Linking Algorithm with the prefrontal–basal ganglia–thalamic circuit,[17] and the Autonomic Stress drive with the amygdala–HPA axis.[13,21] All cited neuroscience studies post-date the XMMM's 2002–2003 original specification,[18,19] making these convergences temporally independent corroborations. A component-by-component mapping with five quantitative predictions is in Supplementary Document S2.

The XMMM occupies a distinct position among existing cognitive architectures (Fig. 4). Against ACT-R and SOAR, it eliminates the module boundary between feeling and thinking. Against Global Workspace Theory, consciousness is located in the continuous competition of drive emotions and sensory states for executive attention. Against the Free Energy Principle, the objective—homeostatic deprivation—is measurable in physiological units. Against standard RL, reward is intrinsic. The XMMM makes a functional completeness claim: every cognitive function performed by an identified brain structure is accounted for by one or more components (Supplementary Document S2, Section 2.7).

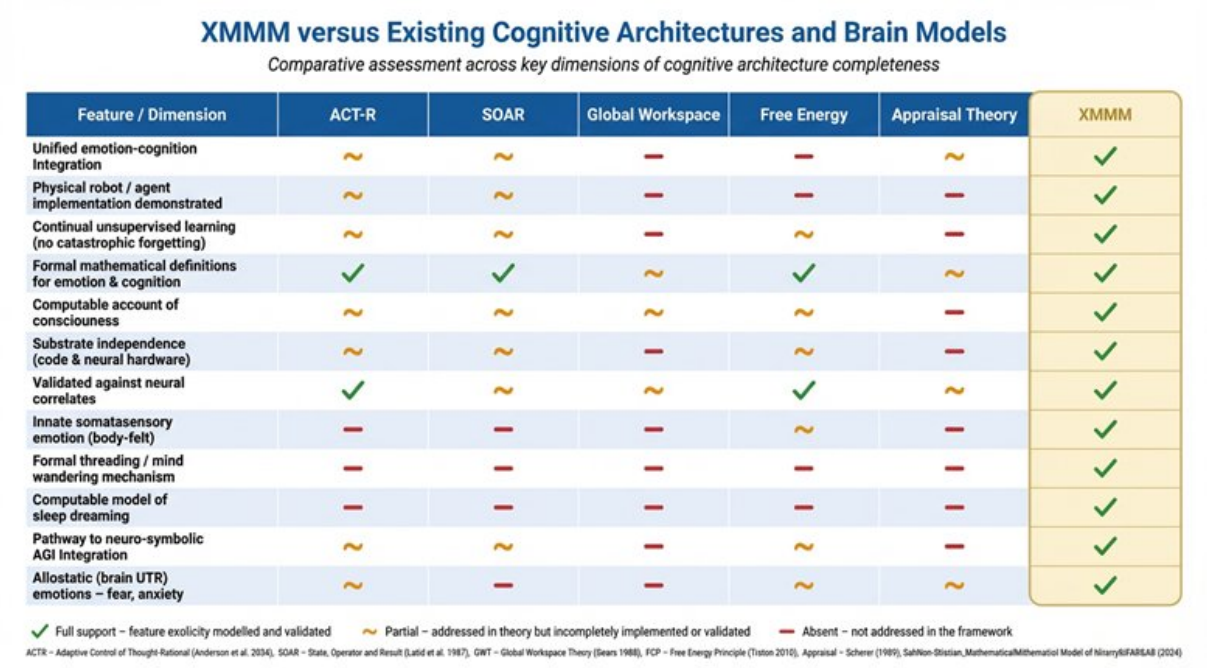

Figure 4 XMMM versus existing cognitive architectures and brain models. Comparative assessment across key dimensions of cognitive architecture completeness. Full support (✓); partial support (∼); absent (–).

The XMMM generates five specific, falsifiable predictions: (1) effective cognitive cycle rate should decrease as $\mathcal{O}(|\mathcal{A}| \log |\mathcal{A}|)$ with associative database size;[16,24] (2) homeostatic drive induction should produce predictable insular cortex activation at specific body locations with magnitudes proportional to $\mathrm{ES}_i(t)$;[13] (3) hippocampal encoding strength should be a monotonic function of emotional intensity at storage time;[25] (4) Default Mode Network (DMN) activity should decrease as a linear function of prime drive ES magnitude, onset at $\theta_{\mathrm{act}} = 0.10$;[26]

and (5) psychiatric conditions correspond to specific measurable derangements of architectural parameters.[27,28] Each prediction carries an explicit falsification condition (Supplementary Document S2, Section 3).

The XMMM provides two separable consciousness claims. Access consciousness is fully computable: a state is access-conscious if and only if it is currently represented in any input stream—exteroceptive sensory, proprioceptive, or Body Map homeostatic— delivered to the Linking Algorithm. Cross-system verification of phenomenal states is physically impossible: a logical consequence of the physical uniqueness of every Body Map and the physical unrealisability of lossless cross-system representational transfer. The Hard Problem is thereby reframed as a derived physical insolubility (Supplementary Document S3, Section 2.2).

Several limitations warrant note. Physical agents used simplified sensory arrays and small association databases; all results derive from observed learning trajectories rather than systematic controlled-trial analyses. Computationally, serial logic-loop retrieval scales as $\mathcal{O}(|\mathcal{A}| \log |\mathcal{A}|)$; Supplementary Document S4 describes the neuro-symbolic hybrid architecture resolving this constraint (Fig. 5).

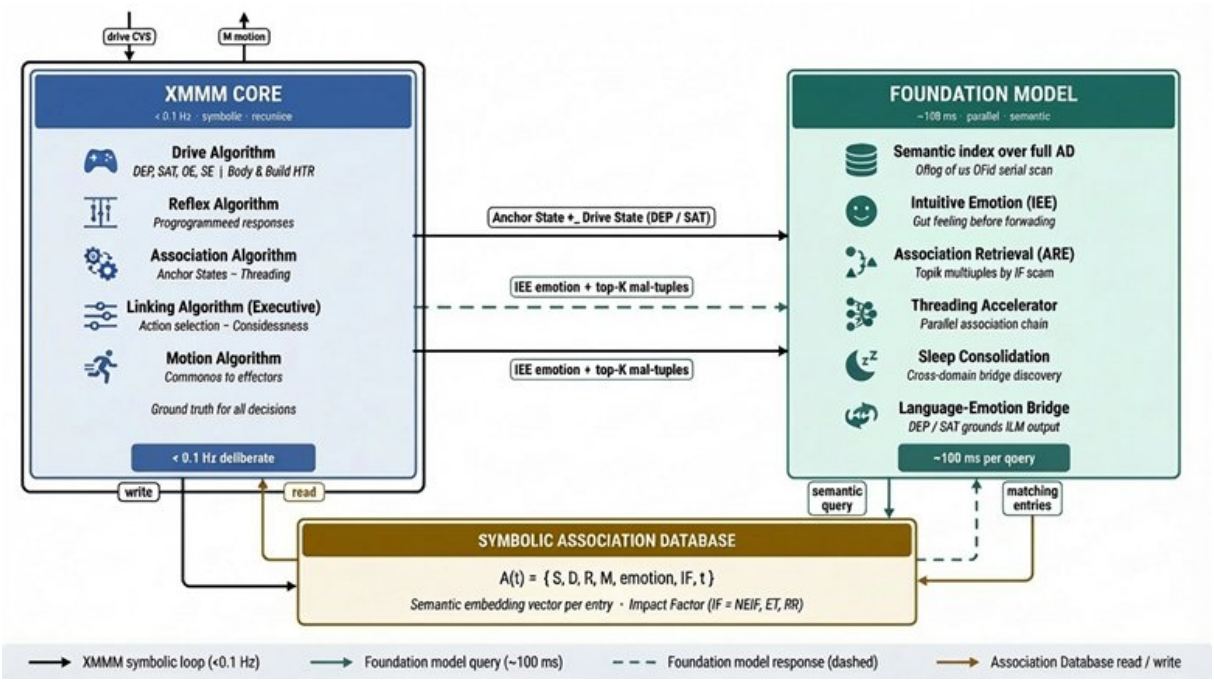

Figure 5 Proposed neuro-symbolic AI architecture based on the XMMM. Four integration planes—Neural Association Index (NAI), Neural Gut-Feel Aggregator (NGA), Neural Threading Accelerator (NTA) and Foundation Model Seeding Interface (FMSI)—resolve the serial XMMM's computational scalability constraint while preserving motivational integrity. The foundation model component is a knowledge source strictly subordinate to the XMMM executive, providing alignment by construction. Full specification in Supplementary Document S4.

The XMMM demonstrates that a unified framework for emotion, cognition and behaviour is achievable from a single control-theoretic axiom. An initiated language acquisition project[22]—modelling vocal outputs as motor skills learned to elicit satiation—will test whether communicative patterns emerge from drive-motivated reinforcement without a dedicated language faculty.

Methods

A. Formal model specification

Architecture overview and substrate independence. The XMMM comprises five substrate-independent functional algorithms integrated by a Linking Algorithm acting as the executive controller. The logic loop executes at approximately 10 Hz in simple implementations. The five algorithms are substrate-independent: the same logical relationships hold whether implemented in C++, Java, Python or future computational substrates.

Variables and notation. All notation follows ISO 80000-2. Single-letter scalars are italic in math mode: t, i, j. Multi-letter variable names are Roman upright: $\mathrm{CV}_i(t)$, SP_i, $\mathrm{ES}_i(t)$, $\mathrm{DE}_i(t)$, $\mathrm{SE}_i(t)$. Subscript running indices are italic; subscript descriptor labels are Roman (e.g., $\mathrm{CV}_{\mathrm{max},i}$). The emotion state vector is bold upright: $\mathbf{D}_j$. Databases and sets use calligraphic font: $\mathcal{A}$.

Autonomic stress drive and allostatic coupling. The Autonomic Stress (AS) drive activates synchronously with every other drive, producing a consolidated valence signal from reflexes, pre-programmed phobias, homeostatic deprivations and allostatic recalls. The Body State Override Reflex (BSOR $\in [0, 1]$) generates an anticipatory autonomic stress modulation consistent with dopaminergic prediction-error signalling.[20]

Association algorithm and Impact Factor. Associations are the fundamental unit of memory. Each association $A_j \in \mathcal{A}$ is a tuple:

$$A_j = \{\mathrm{AS}_j,\ \mathrm{EM}_j,\ \mathbf{D}_j,\ \mathrm{IF}_j,\ t_j\} \tag{7}$$

where AS_j is the Anchor State (compound sensory and drive-emotional context at encoding); EM_j is the executed motor action; $\mathbf{D}_j$ is the emotion state vector; $\mathrm{IF}_j \in (0, 1]$ is the Impact Factor; and t_j is the cycle of experience. The Impact Factor is reinforced at satiation events:

$$\mathrm{IF}_j(t{+}1) = \alpha \cdot \mathrm{IF}_j(t) + (1 - \alpha) \cdot \mathrm{SE}_{\mathrm{PD}}(t), \quad \alpha \approx 0.85\text{–}0.95 \tag{8}$$

and undergoes passive decay otherwise:

$$\mathrm{IF}_j(t{+}1) = \beta \cdot \mathrm{IF}_j(t), \quad \beta \approx 0.97\text{–}0.99 \tag{9}$$

Actions are selected by:

$$\mathrm{EM}^* = \arg\max_j \big(\mathrm{IF}_j \cdot \mathrm{Sim}(\mathrm{AS}(t),\ \mathrm{AS}_j)\big) \tag{10}$$

Relationship to standard reinforcement learning. In standard RL, reward is a scalar signal provided by the task environment. In the XMMM, reward is intrinsic and homeostatic: $\mathrm{SE}_i(t)$ is generated internally from the rate of change of the drive error signal. Numerical benchmarking

against RL baselines is not presented because the two frameworks operate on incommensurable motivational primitives; Fig. 4 constitutes the appropriate comparison.

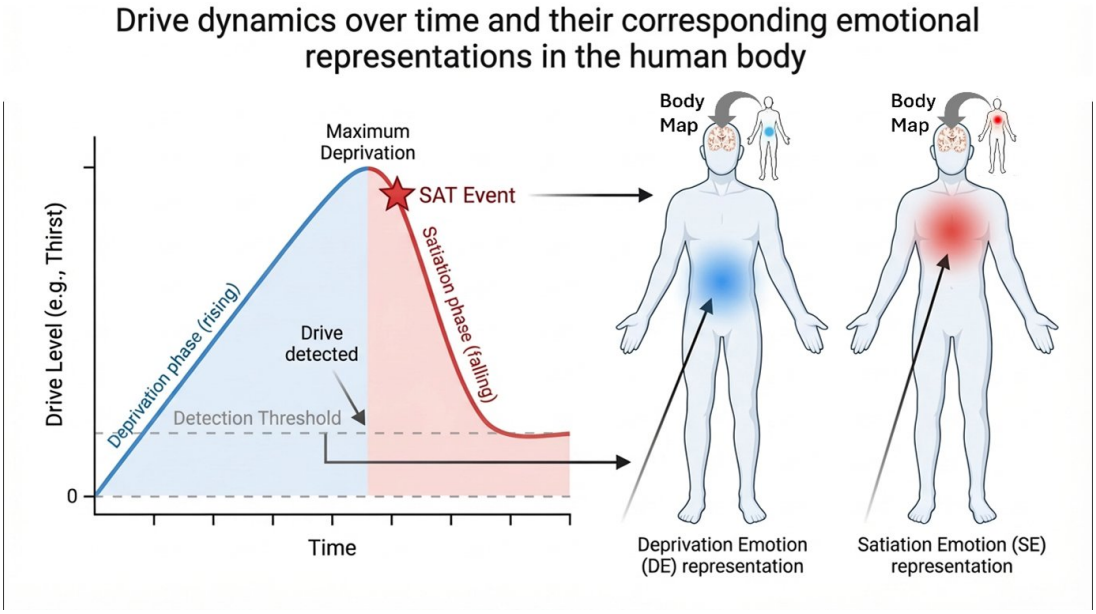

Figure 6 Drive dynamics over time and the corresponding somatosensory emotional representations in the Body Map. The deprivation phase (blue, rising ES_i) generates a body-located aversive Deprivation Emotion (DE; equation 2). At the satiation event (SAT), the drive error begins to fall, generating a Satiation Emotion (SE; equation 3). Both are placed in the Body Map at drive-specific body locations (right panels).

B. Algorithmic implementation

Threading and action retrieval. Threading Focus controls the search mode:

$$\mathrm{Focus}(t) = \mathrm{ES}_{\mathrm{PD}}(t) = |\mathrm{DE}_{\mathrm{PD}}(t)| \ \in [0, 1] \tag{11}$$

$$\mathrm{mode}(t) = \begin{cases} \text{Directed} & \text{if } \mathrm{Focus}(t) \geq \theta_{\mathrm{act}} \\ \text{Undirected} & \text{otherwise} \end{cases} \tag{12}$$

Activation threshold: $\theta_{\mathrm{act}} = 0.10$ (Fig. 7).

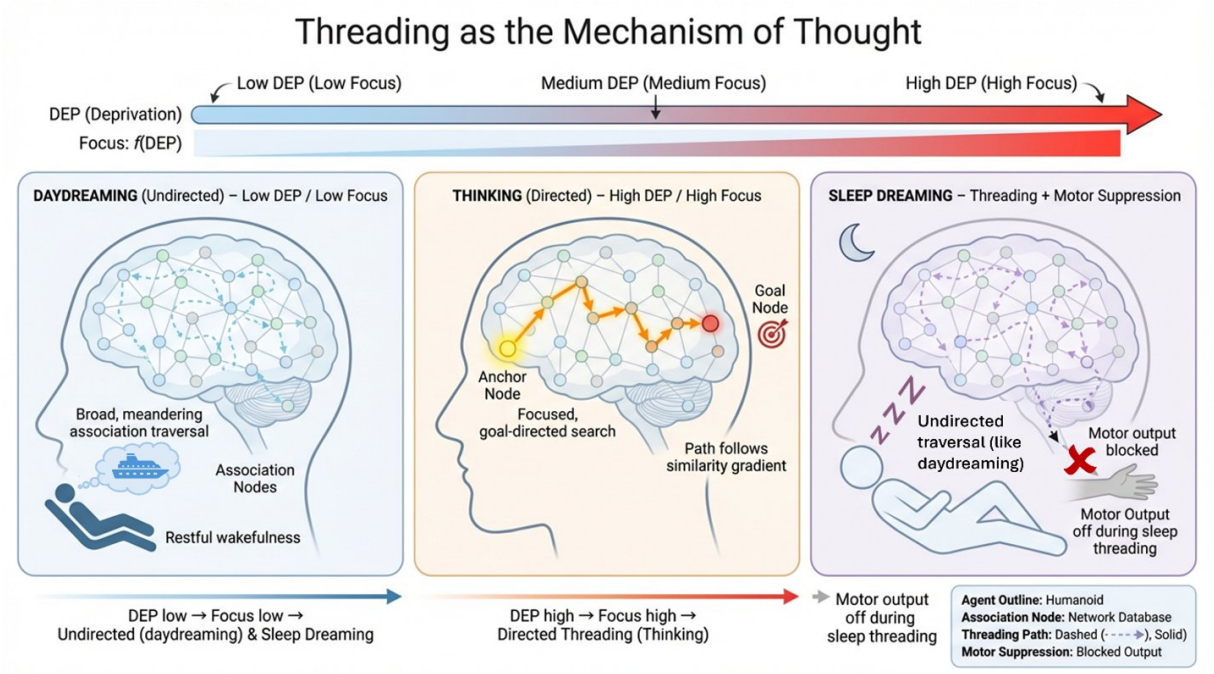

Figure 7 Threading as the mechanism of thought. Schematic showing undirected threading (daydreaming; low deprivation, low Focus), directed threading (thinking; high deprivation, high Focus), and sleep dreaming (threading with motor output suppressed) as modulations of the same underlying association traversal process governed by equations 11–12.

Reward-based backpropagation and its distinction from temporal-difference learning. When a satiation event is detected at cycle t^*, propagation occurs backwards through a Temporal Buffer (TB) of depth $K = 3$–5 cycles (≈300–500 ms), tagging each preceding Anchor State as an approach cue. Unlike temporal-difference (TD) learning, RBP does not maintain a value function; it performs a single-pass causal attribution. TD learning predicts a smooth exponential decay of reinforcement with increasing delay; RBP predicts a step function—testable in conditioning paradigms varying cue-to-satiation onset asynchrony.

Logic loop pseudocode. Each cycle of the XMMM executes the following sequence, where ES_i is the error signal (eq. 1), DE_i the Deprivation Emotion (eq. 2), SE_i the Satiation Emotion (eq. 3), PD the Prime Drive (eq. 4), AS the Autonomic Stress drive, A_j an association tuple (eq. 7), IF_j its Impact Factor (eq. 8), RBP Reward-based Backpropagation, and BSOR the Body State Override Reflex:

```
FOR each cycle t:
  1. Read sensors → update sensory state S(t)
  2. For each drive i:
       compute ES_i(t) = |CV_i(t) - SP_i| / CV_max,i
       determine regime (dES/dt): assign DE_i or SE_i
  3. Compute AS drive from {reflexes, phobias,
                            homeostatic DEs, allostatic recalls}
  4. PD(t) = arg max_i |D_i(t)|
  5. IF |D_PD(t)| > θ_act:
       EM* = arg max_j (IF_j * Sim(AS(t), AS_j))  [directed retrieval]
       Execute EM*
  6. ELSE:
       Sample association by IF-weighted similarity  [undirected threading]
  7. IF satiation event detected: trigger RBP on temporal buffer
  8. Store new association A(t) = {AS(t), EM(t), D(t), IF_init, t}
  9. Update all IFs: reinforcement or passive decay
 10. Linking Algorithm formats EM* as motor command vector:
       MC(t) = {effector_id, direction, magnitude, duration}
       Motion Algorithm executes MC(t); checks Reflex overrides first
 11. GOTO step 1 (recursive loop; continues until agent shutdown)
```

Complexity analysis. Per-cycle: drive update $\mathcal{O}(N)$; retrieval with sorted index $\mathcal{O}(|\mathcal{A}| \log |\mathcal{A}|)$; directed threading at depth k: $\mathcal{O}(k \cdot |\mathcal{A}| \log |\mathcal{A}|)$. The four-plane neuro-symbolic hybrid (Supplementary Document S4) reduces retrieval to $\mathcal{O}(\log |\mathcal{A}|)$.

C. Simulation environments

Project Simmy. Virtual agent in C++ with OpenGL rendering; eleven active drives; Association Database supporting ≈100,000 associations; 10 Hz logic loop. A companion platform (Project Volksbot) implemented the same algorithms on a virtual Fraunhofer Volksbot model.

D. Robot platform

Project Troopy. Java implementation on a wheeled mobile platform (≈1 m^2 arena). Twelve active drives; logic loop 1–3 Hz. The hunger drive proxy ($CV_{hunger}(t)$ rising over ≈4 minutes independently of any sensor) demonstrates the Epistemic Isolation Principle.

E. Experimental procedures

Design rationale and navigation experiments. Trials are not statistically equivalent because each begins from a distinct agent state; results are reported as learning trajectories. Performance metric: time-to-reward (Simmy) and number of movements to reward (Troopy) across successive trials, with no externally specified reward function.

Problem-solving, social bonding and biological validation. Full protocols are provided in Supplementary Documents S1 and S2. The pre-registered social bonding control experiment (two

agent cohorts with reversed panel assignments) is specified in Supplementary Document S1, Section 6.2.

References

1. Craig, A. D. How do you feel—now? *Nat. Rev. Neurosci.* **10**, 59–70 (2009).
2. Friston, K. The free-energy principle. *Nat. Rev. Neurosci.* **11**, 127–138 (2010).
3. Barrett, L. F. & Simmons, W. K. Interoceptive predictions in the brain. *Nat. Rev. Neurosci.* **16**, 419–429 (2015).
4. Anderson, J. R. et al. An integrated theory of the mind. *Psychol. Rev.* **111**, 1036–1060 (2004).
5. Laird, J. E., Newell, A. & Rosenbloom, P. S. SOAR. *Artif. Intell.* **33**, 1–64 (1987).
6. Baars, B. J. *A Cognitive Theory of Consciousness* (Cambridge Univ. Press, 1988).
7. Sternson, S. M. Hypothalamic survival circuits. *Neuron* **77**, 810–824 (2013).
8. Balleine, B. W. & O'Doherty, J. P. Human and rodent homologies in action control. *Neuropsychopharmacology* **35**, 48–69 (2010).
9. Berthoud, H.-R. & Morrison, C. The brain, appetite, and obesity. *Annu. Rev. Psychol.* **59**, 55–92 (2008).
10. Bowlby, J. *Attachment and Loss, Vol. 1* (Basic Books, 1969).
11. Block, N. On a confusion about a function of consciousness. *Behav. Brain Sci.* **18**, 227–247 (1995).
12. Craig, A. D. How do you feel? *Nat. Rev. Neurosci.* **3**, 655–666 (2002).
13. Livneh, Y. & Andermann, M. L. Cellular activity in insular cortex. *Neuron* **109**, 3576–3593 (2021).
14. Zimmerman, C. A. et al. Thirst neurons anticipate homeostatic consequences. *Nature* **537**, 680–684 (2016).
15. Teyler, T. J. & DiScenna, P. The hippocampal memory indexing theory. *Behav. Neurosci.* **100**, 147–154 (1986).
16. Lega, B. C., Jacobs, J. & Kahana, M. Human hippocampal theta oscillations. *Hippocampus* **22**, 748–761 (2012).
17. Haber, S. N. & Knutson, B. The reward circuit. *Neuropsychopharmacology* **35**, 4–26 (2010).
18. Van Schalkwyk, R. A Method and Device to Artificially Reproduce Living Creature Functionality. Patent 2002/1207 (South Africa, 2002).
19. Van Schalkwyk, R. A Method and Device to Illustrate Living Creature Functionality. Patent 2003/0850 (South Africa, 2003).
20. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
21. Cisek, P. & Kalaska, J. F. Neural mechanisms for interacting with a world of action choices. *Annu. Rev. Neurosci.* **33**, 269–298 (2010).

22. Van Schalkwyk, R. & Dehbozorgi, A. Artificial agent language development based on the XMMM. Preprint `https://doi.org/10.13140/RG.2.2.19913.56165` (2024).

23. Scherer, K. R. Appraisal theory. In *Handbook of Cognition and Emotion* 637–663 (Wiley, 1999).

24. Biba, T. M. et al. Episodic memory encoding fluctuates at a theta rhythm. *Nat. Hum. Behav.* (2026). `https://doi.org/10.1038/s41562-026-02416-5`

25. Dolcos, F., LaBar, K. S. & Cabeza, R. Amygdala-hippocampus interaction predicts emotional memory. *Neuron* **42**, 855–863 (2004).

26. Buckner, R. L., Andrews-Hanna, J. R. & Schacter, D. L. The brain's default network. *Ann. N.Y. Acad. Sci.* **1124**, 1–38 (2008).

27. Huys, Q. J. M. et al. Computational psychiatry. *Nat. Neurosci.* **19**, 404–413 (2016).

28. Redish, A. D. Addiction as a computational process gone awry. *Science* **306**, 1944–1947 (2004).

29. Lee, C. R. et al. The neural circuitry of social homeostasis. *Cell* **184**, 1500–1516 (2021).

30. Huang, W.-C. et al. A hypothalamic circuit for social homeostasis. *Nature* (2025). `https://doi.org/10.1038/s41586-025-08653-6`

31. Rappeneau, V. & Castillo Díaz, F. Oxytocin and dopamine signalling convergence. *Neurosci. Biobehav. Rev.* **161**, 105675 (2024).

Acknowledgements

The lead author thanks co-authors D.C. and C.E.A. for neuroscientific consultation, and A.D. for neurolinguistic advice, and acknowledges Prof. Judea Pearl for prompting formal documentation of the algorithms. AI platform Claude (Anthropic) assisted with manuscript preparation and mathematical formalism review.

Author contributions

R.V.S. developed, implemented and tested the XMMM. D.C. provided neuroscientific and psychopathology consultation. C.E.A. provided neuroscientific consultation, experimentation and manuscript review. A.D. provided neurolinguistic consultation and manuscript review. R.V.S. wrote the manuscript with input from all authors.

Data and code availability

Raw data available from the lead author. Java and C++ source code available on Zenodo at [doi.org to be provided on acceptance].

Competing interests

The authors declare no competing interests. R.V.S. is founder of Xzistor LAB, a non-commercial research initiative, and holds intellectual property for the XMMM.

SUPPLEMENTARY INFORMATION

Supplementary Methods and Supplementary Notes 1–4

to

"A Unified Control-theoretic Cognitive Architecture for Emotion, Cognition and Adaptive Behaviour in Biological and Artificial Agents"

Van Schalkwyk, R. Cook, D. Alvarez, C.E. Dehbozorgi, A.

Xzistor LAB, Bristol, United Kingdom · 2026

rocco.vanschalkwyk@xzistor.com | www.xzistor.com

Contents overview:

Contents

Document S0 — Methods Extended

The Xzistor Mathematical Model of Mind

Formal Specification of Algorithms, Equations and Complexity Analysis

Document S0 (Methods Extended) · Supplementary to the Methods section of *"A unified control-theoretic architecture for emotion, cognition and adaptive behaviour in biological and artificial agents"* · Van Schalkwyk, R., Cook, D., Alvarez, C.E., Dehbozorgi, A. · Xzistor LAB, Bristol (UK) · 2026

This document (S0) provides the complete formal mathematical specification of the XMMM: notation conventions, equation derivations, pseudocode, complexity analysis, and summary equation table. It is the canonical reference for all mathematical expressions cited in the main manuscript.

System Architecture and the Cyclic Logic Loop

The XMMM is a deterministic, real-time cognitive architecture that models the biological brain as a closed-loop, multi-variable adaptive control system. The architecture serialises brain function into a cyclic logic loop executed at approximately 10 Hz (one complete cycle ≈ 100 ms), consistent with the dominant timescale of conscious deliberation in vertebrates. Each cycle reads environmental sensors, evaluates internal homeostatic state, generates emotional representations, searches associative memory, selects an effector action, and updates stored associations before repeating.

The architecture comprises five algorithmic building blocks — the Sensing, Drive, Reflex, Association, and Motion Algorithms — integrated by a sixth master routine, the Linking Algorithm (Fig. 1). The five building blocks are substrate-independent functional modules: they may be implemented as compiled procedural code (C++, Java, Python), as software objects in an agent simulation, or as firmware in an embedded robotic controller, or in emerging substrate technologies such as photonic, quantum, or biological processors, without altering the logical relationships described below.

Notation

Single-letter variables (for example α, β, t, i, j, N, K) are typeset in italic. Multi-letter variable names (for example CV, ES, DE, SE, IF, BSOR) are typeset in Roman (upright), consistent with ISO 80000-2 and *Nature* journal style. Bold upright lowercase denotes vectors (e.g. $\mathbf{w}$). Uppercase script ($\mathcal{A}$, $\mathcal{D}$) denotes sets or databases. Subscripts and superscripts that are running indices (i, j, k) are italic; those that are descriptive labels (for example max, thirst, pred, obs, act) are Roman. The symbol t denotes discrete time, counted in logic-loop cycles (1 cycle = 100 ms unless otherwise specified). The operator $|\cdot|$ denotes absolute value and $\mathrm{d}x/\mathrm{d}t$ denotes the finite-difference derivative of x over one cycle.

Core Symbol Glossary

Symbol	Name	Definition / Range
$\mathrm{CV}_i(t)$	Control Variable	Physiological sensor signal for drive i
SP_i	Set-point	Homeostatic equilibrium value for drive i
$\mathrm{ES}_i(t)$	Error Signal	Normalised deviation $\lvert\mathrm{CV}_i - \mathrm{SP}_i\rvert/\text{range} \in [0, 1]$
$D_i(t)$	Drive value	Signed drive strength $\in [-1, +1]$
$\mathrm{DE}_i(t)$	Deprivation Emotion	Aversive pseudo-somatosensory signal $\in [-1, 0]$
$\mathrm{SE}_i(t)$	Satiation Emotion	Appetitive pseudo-somatosensory signal $\in [0, +1]$
$\mathrm{AS}_i(t)$	Autonomic Stress (drive i)	Allostatic coupling output $\in [-1, +1]$
$\mathrm{AS}(t)$	Net Autonomic Stress	Weighted aggregate over all drives $\in [-1, +1]$
$\mathrm{PD}(t)$	Prime Drive	Most urgent drive at cycle t
IF_j	Impact Factor	Reinforcement weight of association $j \in [0, 1]$
A_j	Association	Stored snapshot: $\{\mathrm{AS}_j, \mathrm{EM}_j, \mathbf{D}_j, \mathrm{IF}_j, t_j\}$
$\mathrm{TB}(t)$	Temporal Buffer	Ring buffer of K preceding associations
$\mathrm{BSOR}(t)$	Body State Override Reflex	Limbic enhancement scalar $\in [0, 1]$
K	Buffer depth	Number of cycles in temporal buffer (default 3–5)
N	Drive count	Total number of active drives

Drive Algorithm: Homeostatic and Allostatic Control Loops

The Drive Algorithm models each physiological or allostatic need as an independent negative-feedback control loop. The key insight of the XMMM is that raw error signals from these loops are never presented directly to the executive (cognition). Instead, they are first transformed into pseudo-somatosensory emotional representations — the sole motivational channel available to the executive — as described in M3.

Error Signal Computation

For drive i, the normalised error signal $\mathrm{ES}_i(t)$ measures the fractional deviation of the relevant control variable $\mathrm{CV}_i(t)$ from its homeostatic set-point SP_i:

$$\mathrm{ES}_i(t) = \frac{|\mathrm{CV}_i(t) - \mathrm{SP}_i|}{\mathrm{CV}_{\mathrm{max},i}} \in [0, 1] \tag{1}$$

where $\mathrm{CV}_{\mathrm{max},i}$ is the maximum physiological excursion of control variable i used for normalisation. $\mathrm{ES}_i = 0$ indicates perfect homeostasis; $\mathrm{ES}_i = 1$ indicates a maximally critical deficit. For multi-variable drives (such as thirst, which integrates plasma osmolality O, blood volume V, and angiotensin II level A_{II}), the consolidated error signal can be modelled as a weighted linear

combination:

$$\mathrm{ES}_{\mathrm{thirst}}(t) = w_O \cdot O(t) + w_V \cdot V(t) + w_{\mathrm{AII}} \cdot A_{\mathrm{II}}(t) \tag{2}$$

with weights $w_O + w_V + w_{\mathrm{AII}} = 1$ established from neurophysiological data. The same weighted-sum principle applies to any drive with multiple control variables.

Homeostatic versus Allostatic Drives

The XMMM partitions drives into two functional classes according to how their error signals can be modified.

Homeostatic drives admit change in ES_i only through physical changes in their control variables. Examples include thirst, hunger, cold, warm, pain (localised), fatigue, and urge to urinate. Once the organism is removed from the physiological perturbation, the error signal recovers automatically.

Allostatic drives admit change in ES_i both from physical control-variable changes and from recalled associative memories. The biological autonomic stress response (HPA axis, fight-or-flight), anger, acute fear, nausea, sexual arousal, and bonding are allostatic. The key property of allostatic drives is that aversive anticipation — the cognitive recall of a past deprivation event — is sufficient to raise the error signal in the absence of any current physiological threat. This allows the agent to act pre-emptively out of fear, anxiety, or anticipatory pleasure.

The Autonomic Stress Drive and Inter-Drive Coupling

The autonomic stress drive occupies a unique architectural position: it is activated in parallel with every other homeostatic and allostatic drive whenever that drive enters deprivation. The coupling is modelled as:

$$\mathrm{ES}_{\mathrm{AS}}(t) = \max_i \left[w_i \cdot \mathrm{ES}_i(t) \right] + \sum_{j \in \mathrm{MEM}} \left[\mathrm{IF}_j \cdot \mathrm{AS}_{\mathrm{recalled},j}(t) \right] \tag{3}$$

where the first term captures the maximum instantaneous drive stress and the second term accumulates allostatic contributions from recalled associations (memories) weighted by their impact factors IF_j. This coupling is the mechanism by which the XMMM models the fact that environmental cues associated with past deprivation events can themselves provoke autonomic stress, even in the absence of current physiological need. It should be noted that AS re-evocation is not the sole consequence of association retrieval: all allostatic emotional components stored within the recalled association — including nausea, anger, sexual arousal, and bonding responses — are simultaneously re-evoked, producing the compound emotional response characteristic of the original event (see M6).

Inter-Drive Suppressive Coupling

In biological organisms, concurrent activation of multiple homeostatic drives does not produce a simple summation of their individual urgencies. Specific cross-inhibitory relationships between drives have been documented: in particular, dehydration suppresses appetite through hy-

pothalamic cross-inhibition, conserving metabolic resources during water deficit. The XMMM implements this through a fractional inhibitory weighting applied between selected drive pairs:

$$\mathrm{ES}_i(t) \;\leftarrow\; \mathrm{ES}_i(t) \cdot \left[1 - \alpha_{ij} \cdot \mathrm{ES}_j(t)\right] \tag{M2.5-i}$$

where $\alpha_{ij} \in [0, 1]$ is the suppressive coupling coefficient from drive j onto drive i. In the reference implementation, the thirst-to-hunger suppressive coefficient $\alpha_{\text{thirst}\to\text{hunger}} = 0.10$, consistent with the known fractional magnitude of hypothalamic cross-inhibition reported in the mammalian literature. Under simultaneous hunger and thirst deprivation, the agent's effective hunger urgency is therefore reduced by 10% of the current thirst error signal, producing drive prioritisation under combined conditions that more closely resembles mammalian homeostatic trade-offs than the non-suppressive case. This constraint was incorporated following observation of implausible concurrent-drive behaviour in early Troopy implementations (see Document S1, Section 1.5.1), and its convergence with the known biological mechanism constitutes an independent validation of the Drive Algorithm's control-theoretic structure.

Emotion Generation: Deprivation and Satiation Representations

The conceptual core of the XMMM is the claim that the executive part of the modelled brain is never granted direct access to raw error signal magnitudes. Instead, all drive information is first converted into pseudo-somatosensory emotional representations — numerical quantities that mimic the felt, body-located quality of biological emotions — and placed in the Body Map. The executive reads the Body Map, not the drives. This architectural choice has three critical consequences: (i) it creates a single unified motivational language for the executive across all drive types; (ii) it embeds drive information in a spatial body representation, generating embodied emotional awareness; and (iii) it limits the information channel between physiological state and cognition to precisely what the agent "feels", consistent with human introspective reports.

Deprivation Emotion

During the deprivation regime — defined as the interval when $\mathrm{ES}_i(t) > 0$ and is not decreasing — the drive generates a Deprivation Emotion (DE). The DE representation is a negative scalar in $[-1, 0]$, where the magnitude encodes the urgency of the unmet need:

$$\mathrm{DE}_i(t) = -\mathrm{ES}_i(t) \;\in [-1, 0] \tag{6}$$

The negative sign is a formal attribute that signals to the executive that the current state requires corrective action (approach or avoidance). A DE of -1 represents maximal deprivation; a DE of 0 represents homeostasis. The DE is placed as a pseudo-somatosensory entry in the Body Map with attributes encoding: the identity of drive i, the spatial body location associated with that drive, the sign (deprivation), and the magnitude.

Satiation Emotion

The Satiation Emotion (SE) is generated during the satiation regime — defined as the interval when $\mathrm{ES}_i(t)$ is actively decreasing (the derivative is negative). The SE encodes the rate of homeostatic recovery, normalised to the maximum physiologically achievable rate:

$$\mathrm{SE}_i(t) = \min\left(\frac{|\mathrm{dES}_i(t)/\mathrm{d}t|}{\dot{E}S_{\max,i}},\ 1.0\right) \in [0, +1] \tag{7}$$

where $\dot{E}S_{\max,i}$ is the maximum possible rate of error signal reduction for drive i (e.g., determined by the fastest possible fluid absorption rate for thirst). A SE of $+1$ represents instantaneous complete recovery — the fastest possible quench. A SE of $+0.01$ represents the slowest perceptible improvement. The positive sign signals approach, reinforcement, and subjective pleasure.

The interpretation is significant: biological pleasure and reward are not directly proportional to the current state of need but to the *speed of recovery* from need. A person who has been profoundly dehydrated does not subjectively feel more pleased by each mouthful of water than a mildly thirsty person — rather, the rate of quenching, relative to baseline rate, determines the magnitude of the satiation emotion. This is consistent with the Weber–Fechner and derivative-based accounts of sensory coding.

Signed Drive Value and the Deprivation–Satiation Transition

For each drive, the model maintains a signed drive value $D_i(t) \in [-1, +1]$ that summarises the current emotional state as a single scalar:

$$D_i(t) = \begin{cases} \mathrm{DE}_i(t) & \text{if } \mathrm{dES}_i/\mathrm{d}t \geq 0 \quad \text{(deprivation regime)} \\ \mathrm{SE}_i(t) & \text{if } \mathrm{dES}_i/\mathrm{d}t < 0 \quad \text{(satiation regime)} \end{cases} \tag{8}$$

The transition from $D_i < 0$ to $D_i > 0$ at the satiation event — the moment $\mathrm{dES}_i/\mathrm{d}t$ first becomes negative — is the primary reinforcement trigger in the architecture. It is at this transition that the temporal buffer is read and reward is backpropagated to preceding associations (see M6.2).

Body Map Placement and Pseudo-Somatosensory Representation

DE and SE representations are not abstract numerical variables: they are placed in a Body Map data structure with explicit body-location attributes. For example, the thirst drive's DE is associated with the epigastric and throat regions; the thirst SE with the same area plus an oral-pharyngeal component. Pain drives are each associated with a specific body-region identifier matching the nociceptive input location.

The pseudo-somatosensory quality arises because these emotion representations are indistinguishable in format to the executive from genuine exteroceptive and interoceptive sensory signals arriving from the same body locations. From the executive's perspective, hunger "feels" like it comes from the gut because the hunger DE is mapped to the gut region of the Body Map — not because there is a dedicated "hunger feeling" structure. This architecture directly implements

Damasio's (1994, 2018) as-if body loop at a mechanistic, computable level.

The Body Map need not be a topographic physical array; in digital implementations it is a relational database table where each row is a sensor or emotion representation and columns store: `region_id`, `modality_id`, `signal_strength`, `signed_value`, `body_location_x/y/z`, and a flag distinguishing genuine sensory from pseudo-somatosensory entries. The Linking Algorithm reads all rows of this table without needing to distinguish their origin.

Combinatorial Emotion Space

With N distinct drives, each capable of DE states at resolution r_D and SE states at resolution r_S, the total number of simultaneous emotional configurations accessible to the agent is:

$$\Omega = (r_D \times r_S)^N \tag{9}$$

For the XMMM reference implementation with $N = 20$ drives, $r_D = r_S = 100$ (1% increment resolution):

$$\Omega = (100 \times 100)^{20} = 10^{42} \tag{10}$$

This exceeds the estimated number of grains of sand on Earth ($\approx 7.5 \times 10^{18}$) by 23 orders of magnitude. The combinatorial richness of the emotion space arises not from coding individual emotions but from the intersection of twenty simultaneously active drives, each continuously valued — a consequence of the control-theoretic architecture. Cultural and social contextualisation of emotion (the learned association of these combinatorial states with specific environmental cues and social contexts) adds a further combinatorial multiplier of effectively unlimited magnitude.

Reflex Algorithm

Reflexes are preprogrammed, non-learned motor responses triggered by specific sensory states or drive states that exceed threshold values. They execute within the same logic-loop cycle as their trigger, bypassing the Association Algorithm and the Linking Algorithm's deliberative phase.

When triggered by a sensory state $S_k(t)$ exceeding a threshold θ_k:

$$R(t) = f_R(S_k(t)) \quad \text{if } S_k(t) \geq \theta_k \tag{11}$$

When triggered by a drive $D_i(t)$ exceeding a threshold θ_i:

$$R(t) = g_R(D_i(t)) \quad \text{if } |D_i(t)| \geq \theta_i \tag{12}$$

Reflexes are classified across two valence conditions and two modifiability classes, giving four types: (1) Involuntary Deprivation Reflex (e.g., shivering, pupil dilation); (2) Learn-modifiable Deprivation Reflex (e.g., withdrawal from a pain source that becomes refined through experience); (3) Involuntary Satiation Reflex (e.g., swallowing, smooth muscle peristalsis); (4) Learn-modifiable Satiation Reflex (e.g., suckling pattern in neonates). Learn-modifiable reflexes have their initial responses overwritten in the Association Database through operant conditioning.

Reflexes from anger (Eq. 4) and acute fear (Eq. 5) are architecturally important: anger reflexes are initialised as forward-approach motor programs directed at the recognised stressor; acute fear reflexes are initialised as withdrawal motor programs. Both are subsequently refined by experience.

Association Algorithm: Operant Learning and Associative Memory

Associations are the fundamental unit of memory in the XMMM. At each logic-loop cycle, the current state of the agent — comprising sensory representations, drive emotion vectors, autonomic stress level, and current effector motions — is encoded as an association snapshot and either stored as a new entry or used to update an existing entry in the Association Database ($\mathcal{A}$).

Association Structure

Each association $A_j \in \mathcal{A}$ is a tuple:

$$A_j = \{\, \mathrm{AS}_j,\ \mathrm{EM}_j,\ \mathbf{D}_j,\ \mathrm{IF}_j,\ t_j \,\} \tag{13}$$

where AS_j (Anchor State) is the compound key — the set of sensory representations S and drive representations D active at the time of encoding; EM_j is the effector motion (action) being performed; $\mathcal{D}_j$ is the full emotion vector $[\mathrm{DE}_1, \mathrm{SE}_1, \mathrm{DE}_2, \mathrm{SE}_2, \ldots, \mathrm{DE}_N, \mathrm{SE}_N]$; AS_j is the autonomic stress scalar; $\mathrm{IF}_j \in [0, 1]$ is the Impact Factor; and t_j is the creation timestamp.

Impact Factor and Reinforcement Update

The Impact Factor (IF) quantifies the learned motivational salience of an association — equivalently, the strength of the reinforcement link between the stored context and the stored action. It is updated at every satiation event according to:

$$\mathrm{IF}_j(t+1) = \alpha \cdot \mathrm{IF}_j(t) + (1 - \alpha) \cdot \mathrm{SE}_{\mathrm{PD}}(t) \tag{14}$$

where $\alpha \in (0, 1)$ is a decay-modulated learning rate (typically 0.85–0.95 in validated implementations) and $\mathrm{SE}_{\mathrm{PD}}(t)$ is the satiation emotion value of the prime drive at cycle t. Equation 14 is a running exponential average: associations that have repeatedly preceded strong satiation events accumulate high impact factors; associations that have not led to satiation decay toward zero. This is the formal equivalent of Thorndike's Law of Effect expressed as a continuous scalar.

When a satiation event does not occur (ES_i remains constant or rising), IF_j is gently decremented:

$$\mathrm{IF}_j(t+1) = \beta \cdot \mathrm{IF}_j(t), \quad \beta \in (0,\, \alpha),\ \beta \text{ typically } 0.97\text{–}0.99 \tag{15}$$

This slow passive decay models the neurophysiological observation that memories that are never reactivated weaken over time, providing a computational correlate of forgetting.

Association Retrieval and Contextual vs. Non-Contextual Recall

At each cycle, the Linking Algorithm queries $\mathcal{A}$ via two search passes, executed sequentially.

The Anchor State AS_j is the compound associative key that uniquely identifies the contextual conditions under which association A_j was formed. It is a structured tuple encoding the sensory state $S(t)$ active at storage time — comprising all active perceptual channel values (visual, auditory, olfactory, tactile, proprioceptive) above their detection thresholds — together with the signed drive emotion vector $\mathbf{D}(t) = [D_1(t), D_2(t), \ldots, D_N(t)]$ encoding the motivational state at storage time. Formally: $\text{AS}_j = \{S_j, \mathbf{D}_j\}$ where S_j is the sensory component and $\mathbf{D}_j$ is the drive emotion component. During retrieval, the similarity measure $\text{sim}(\text{AS}_j, \text{AS}(t))$ operates over this compound key: a retrieved association is contextually appropriate not merely if the sensory context resembles the stored one, but if the motivational state at retrieval also resembles the motivational state at storage. This dual-component matching is what distinguishes contextualised from non-contextualised recall and ensures that the same sensory cue retrieves different actions depending on the current drive state — the sight of a door retrieves different behaviour in a hungry, a thirsty, or a satiated agent. The Anchor State is the direct algorithmic implementation of Teyler and DiScenna's (1986) hippocampal memory indexing theory: a partial cue activates the stored index and re-evokes the full association content (see Document S2, Section 2.1).

Non-contextualised recall: the Association Database is searched by sensory state alone (S-matching), without regard to current drive state. All matching associations contribute their stored autonomic stress values to the running gut-feel estimate — a rapid, non-deliberative first-impression response that contributes directly to the Body State Override Reflex (M7).

Contextualised recall: the Association Database is searched by the full Anchor State — the combination of current sensory state $S(t)$, current prime drive $\text{PD}(t)$, and a tolerance radius ε around the current emotion vector. The best-matching subset $\mathcal{A}_{\text{match}} \subset \mathcal{A}$ is ranked by IF and the highest-IF entry provides the candidate action for the current cycle.

The combined search determines the candidate action $\text{EM}_{\text{candidate}}(t)$:

$$\text{EM}_{\text{candidate}}(t) = \underset{A_j \in \mathcal{A}_{\text{match}}}{\arg\max} \left[\text{IF}_j \cdot \text{sim}(\text{AS}_j, \text{AS}(t))\right] \tag{16}$$

where $\text{sim}(\cdot, \cdot)$ is a normalised similarity measure over the anchor state (inner product for numerical representations, equality for categorical identifiers).

Correlation Sensitivity: Partial-Match Retrieval Threshold

The similarity measure $\text{sim}(\text{AS}_j, \text{AS}(t))$ in Eq. 16 produces a continuous value in $[0, 1]$, but action retrieval is only triggered when this similarity exceeds a threshold parameter termed the Correlation Sensitivity (CS). The CS governs the minimum degree of contextual overlap required

before an association becomes eligible for retrieval:

$$\mathrm{EM}_{\mathrm{candidate}}(t) = \underset{A_j \in \bar{\mathcal{A}}_{\mathrm{match}}}{\arg\max} \left[\mathrm{IF}_j \cdot \mathrm{sim}(\mathrm{AS}_j,\, \mathrm{AS}(t))\right] \quad \text{subject to } \mathrm{sim}(\mathrm{AS}_j,\, \mathrm{AS}(t)) \geq \mathrm{CS}_j \tag{M5.4-i}$$

The default value $\mathrm{CS}_j = 0.10$ (10% overlap threshold) permits broad partial-match retrieval, enabling generalisation to novel sensory contexts. CS is not fixed: it is updated by RBP in response to retrieval outcomes, implementing experience-driven calibration of retrieval precision.

Tightening rule

When a retrieved partial-match association successfully leads to satiation, CS_j is increased by $\gamma_+ = 0.05$, narrowing future retrieval to contexts more similar to the current one. High-stakes associations (e.g., those associated with pain avoidance or maximum hunger satiation) accumulate tight CS values, reducing the risk of inappropriate generalisation.

Relaxation rule

When a retrieved partial-match fails to produce satiation within the temporal buffer window, CS_j is decreased by $\gamma_- = 0.02$, broadening future retrieval radius. This relaxation under failure is the mechanism by which the agent progressively widens its search when immediate attempts are unsuccessful, implementing a systematic broadening of inductive inference under increasing drive pressure.

Drive-pressure global modulation

In addition to association-specific CS update, the prime drive deprivation magnitude applies a global floor reduction to all CS values: $\mathrm{CS}_{\mathrm{eff}} = \mathrm{CS}_j \cdot (1 - \alpha_c \cdot |\mathrm{DE}_{\mathrm{PD}}|)$, where $\alpha_c \approx 0.5$ in the reference implementation. Under high deprivation, the effective CS approaches zero, permitting retrieval of any association with even marginal contextual overlap. This is the formal mechanism of "frantic" behaviour under extreme drive pressure.

The Threading Cache

During directed threading (M8.3), the Linking Algorithm does not immediately commit to a single retrieved association. Instead, it populates a short-term Threading Cache (TC): a ranked list of the highest-IF candidate associations satisfying the current CS threshold, retained across multiple logic-loop cycles to support sequential associative traversal.

The TC stores the top-K candidates from Eq. 16 (typically $K = 3$–5 associations), ordered by $\mathrm{IF}_j \cdot \mathrm{sim}(\mathrm{AS}_j, \mathrm{AS}(t))$. The highest-ranked candidate is moved to the Executive Control Buffer (ECB) — the active working context slot held by the Linking Algorithm — where its visual and emotional content is re-evoked for the executive controller: the agent effectively "imagines" the stored sensory state and its associated emotions before committing to the corresponding motor action. If the imagined state fails to produce a viable approach to satiation — indicated by the absence of a BSOR_+ response to the anticipated trajectory — the next candidate in the TC is retrieved and evaluated in turn.

The sequential evaluation of cached candidates naturally organises along reward gradients. Associations encoding landmark states closer to the reward source carry higher IF values (from more frequent RBP reinforcement) and therefore appear earlier in the cache. When the agent traverses the TC sequentially, it constructs a mentally simulated path from its current location to the reward source — re-evoking landmark cues and associated emotions in the order in which they were originally encountered along successful navigation routes. This is the mechanism by which Troopy was observed to "imagine" routes to food sources during periods of high hunger deprivation. The behaviour requires no route-planning algorithm: it emerges from the IF-weighted sequential structure of cached associations and the emotional re-evocation of each cached state.

The TC is cleared at the satiation event (M8.4) and re-populated at each directed threading cycle. Its maximum depth K_{TC} is a configurable parameter. Note: the symbol K_{TC} (Threading Cache depth) is distinct from K_{TB} (Temporal Buffer depth, §6.1), which governs the retrospective RBP credit-attribution window ($K_{\mathrm{TB}} = 3$–5 cycles).

Reward-based Backpropagation and the Temporal Buffer

Reward-based backpropagation (RBP) is the mechanism by which satiation events retrospectively reinforce not only the action that directly triggered satiation, but also the chain of preceding environmental cues and actions that led to it. This is the XMMM's account of how distant navigational cues (e.g., a green door three corridors from a food source) become secondary reinforcers, enabling agents to navigate to reward sources from increasing distances over repeated trials.

The Temporal Buffer

A temporal buffer TB of depth K stores the K most recent association snapshots prior to the current cycle, as a first-in-first-out ring structure:

$$\mathrm{TB}(t) = \{A(t-K),\ A(t-K+1),\ \ldots,\ A(t-1)\} \tag{17}$$

Default buffer depth $K = 3$–5 cycles (300–500 ms at 10 Hz). This interval is chosen to span the typical approach time from a proximal environmental cue to the reward source, consistent with the timescale of classical conditioning windows reported in rodent and primate learning studies.

Satiation Event Detection

A satiation event for drive i at cycle t is detected when:

$$\frac{\mathrm{dES}_i(t)}{\mathrm{d}t} < 0 \quad \text{AND} \quad \frac{\mathrm{dES}_i(t-1)}{\mathrm{d}t} \geq 0 \tag{28}$$

That is, the error signal first crosses from non-decreasing to decreasing. Detection triggers: (i) the temporal buffer read and RBP update (M6.3); (ii) BSOR prediction-error computation (M7); (iii) IF updates for all associations in the temporal buffer (Eq. 19); (iv) Body Map

replacement of DE_i with SE_i for drive i.

Backpropagation at the Satiation Event

When a satiation event is detected at cycle t^* (that is, $\mathrm{dES_{PD}}/\mathrm{d}t$ first becomes negative), the temporal buffer is read and the IF of each buffered association is updated with a weight that decays exponentially with temporal distance from the satiation event:

$$\Delta\mathrm{IF}_j = \gamma^{(t^* - t_j)} \cdot \mathrm{SE_{PD}}(t^*) \quad \text{for } A_j \in \mathrm{TB}(t^*) \tag{18}$$

where $\gamma \in (0, 1)$ is the temporal discount factor (typically 0.85–0.95), t_j is the cycle index when association A_j was generated, and $\mathrm{SE_{PD}}(t^*)$ is the satiation emotion strength at the satiation event. The earlier the association in the buffer (larger $t^* - t_j$), the smaller the reinforcement it receives. The combined update rule for buffered associations is:

$$\mathrm{IF}_j(t^* + 1) = \mathrm{IF}_j(t^*) + \Delta\mathrm{IF}_j \cdot \big(1 - \mathrm{IF}_j(t^*)\big) \quad \text{[clipped to } [0, 1]] \tag{19}$$

This clipping ensures IF remains a proper probability-like weight in $[0, 1]$. The $(1 - \mathrm{IF}_j)$ multiplier implements a diminishing-returns correction analogous to the Rescorla–Wagner model, preventing already-saturated associations from receiving wasteful additional reinforcement.

The biological substrate of this mechanism is assumed to be the dopaminergic prediction-error signal described by Schultz, Dayan & Montague (1997), combined with the cholinergic modulation of synaptic plasticity that gates Hebbian learning windows. The XMMM does not model individual neurons but provides the system-level functional equivalent of this well-documented neurocomputational process.

Body State Override Reflex: Modelling the Limbic System

The Body State Override Reflex (BSOR) is the XMMM's functional model of limbic system enhancement. It is activated by a discontinuous change in the autonomic stress signal — specifically, when the actual change in AS exceeds the predicted (anticipated) change by a signed threshold. The BSOR temporarily overrides the error signals of multiple drives, creating an artificially elevated emotional state (positive or negative) that both enhances the salience of the immediately preceding and succeeding associations and drives heightened motor responding.

Prediction and Prediction Error

The anticipated autonomic stress change $\Delta\mathrm{AS_{pred}}(t)$ is maintained as a running exponential average of recent observed changes:

$$\Delta\mathrm{AS_{pred}}(t) = \eta \cdot \Delta\mathrm{AS_{pred}}(t-1) + (1 - \eta) \cdot \Delta\mathrm{AS_{obs}}(t-1), \quad \eta \approx 0.90 \tag{20}$$

The autonomic stress prediction error $\delta(t)$ is:

$$\delta(t) = \Delta\mathrm{AS_{obs}}(t) - \Delta\mathrm{AS_{pred}}(t) \tag{21}$$

When $\delta(t)$ exceeds a threshold θ_{BSOR} in magnitude, the BSOR activates:

$$\text{BSOR}(t) = \min\left(\frac{|\delta(t)|}{\delta_{\max}},\ 1.0\right) \cdot \mathbf{1}\big[|\delta(t)| \geq \theta_{\text{BSOR}}\big] \tag{22}$$

Positive and Negative BSOR

The sign of $\delta(t)$ determines the direction of the override. A positive prediction error (better-than-expected reward) produces a BSOR_+ that temporarily enhances satiation emotions:

$$\text{SE}_i^{\text{BSOR}}(t) = \min\big(\text{SE}_i(t) + \text{BSOR}(t) \cdot \kappa_+,\ 1.0\big) \tag{23}$$

A negative prediction error (worse-than-expected or absent reward) produces a BSOR_- that temporarily enhances deprivation emotions:

$$\text{DE}_i^{\text{BSOR}}(t) = \max\big(\text{DE}_i(t) - \text{BSOR}(t) \cdot \kappa_-,\ -1.0\big) \tag{24}$$

where κ_+ and $\kappa_- \in (0, 1]$ are gain parameters (set to 0.6 in reference implementations). The BSOR_--enhanced emotions replace DE and SE in the Body Map for the current and the following τ_{BSOR} cycles (typically $\tau_{\text{BSOR}} = 3$–5 cycles, i.e., 300–500 ms), after which the Body Map reverts to the non-enhanced values.

BSOR reinforcement updates: during the τ_{BSOR} window, the impact factor update (Eq. 14) uses the BSOR_--enhanced SE value, not the baseline SE, thereby amplifying the reinforcement of all associations formed during this window. This provides the model's account of why unexpected rewards are more powerfully reinforcing than expected ones — directly parallel to the prediction-error coding established in the dopaminergic system.

Habituation

With repeated exposure to the same satiation source, $\Delta\text{AS}_{\text{pred}}$ approaches $\Delta\text{AS}_{\text{obs}}$, causing $\delta \to 0$ and therefore $\text{BSOR} \to 0$. The BSOR contribution to reinforcement vanishes upon habituation. This implements the well-known empirical observation that repeated reward exposure reduces the hedonic response (hedonic adaptation) without altering the homeostatic value of the reward.

Formally, the habituation time constant τ_{hab} is:

$$\tau_{\text{hab}} \approx \frac{1}{1 - \eta} \text{ cycles} \tag{25}$$

For $\eta = 0.90$, $\tau_{\text{hab}} \approx 10$ cycles $= 1$ second, consistent with rapid short-term habituation to somatosensory stimuli reported in psychophysics.

Linking Algorithm: Prime Drive Selection, Threading, and Executive Control

The Linking Algorithm is the executive controller that integrates all five building-block algorithms into a coherent, goal-directed action cycle. It runs once per logic-loop cycle and performs

the following operations in sequence: (i) Prime Drive selection; (ii) threading mode selection; (iii) candidate action retrieval; (iv) satiation event detection and RBP trigger; (v) Body Map update.

Prime Drive Selection

The Prime Drive $\mathrm{PD}(t)$ is the drive with the highest absolute emotional urgency at cycle t. It is selected across all N drives and both valence directions:

$$\mathrm{PD}(t) = \underset{i \in \{1,\ldots,N\}}{\arg\max} \; |D_i(t)| \tag{26}$$

where $D_i(t)$ is the signed drive value defined in Eq. 8. The PD functions as the dominant motivational context for the current cycle: it determines which drive-specific associations are preferentially retrieved (contextual recall, M5.3) and which satiation events will trigger RBP. Note that the absolute value in Eq. 26 means that an extremely strong satiation emotion can itself become prime — an agent acutely experiencing the pleasure of eating will continue to pursue eating even if another drive has a non-trivial deprivation, because the current satiation rate exceeds the deprivation urgency of all other drives.

Conflict Resolution for Equal-Urgency Drives

When two drives have equal $|D_i|$, a secondary tiebreaking rule applies based on allostatic priority: allostatic drives take precedence over homeostatic drives at equal magnitude, reflecting the observation that psychosocial and autonomic stress tends to be acted upon with greater urgency than equivalent-magnitude homeostatic deficits. If allostatic type does not resolve the tie, the drive with the highest instantaneous rate of change $\mathrm{d}|D_i|/\mathrm{d}t$ is selected (rising urgency takes precedence over a plateau).

Threading: Directed Cognition versus Mind-Wandering

Once the prime drive is selected, the Linking Algorithm determines the cognitive mode — directed threading or default-mode threading — according to a salience threshold θ_{act}:

$$\mathrm{mode}(t) = \begin{cases} \text{Directed} & \text{if } |D_{\mathrm{PD}}(t)| \geq \theta_{\mathrm{act}} \\ \text{Default} & \text{if } |D_{\mathrm{PD}}(t)| < \theta_{\mathrm{act}} \end{cases} \tag{27}$$

In Directed Threading mode, the Association Database is searched using the full contextual query (Eq. 16) and the highest-IF contextual association provides the action candidate. The agent is goal-directed.

In Default-mode Threading ($\theta_{\mathrm{act}} = 0.10$ in reference implementations), no single drive dominates. The Association Database is searched using only sensory state $S(t)$ without drive anchoring. The returned associations are sampled from a probability distribution weighted by IF values, rather than selecting the maximum-IF entry. This non-deterministic, drive-unanchored search is the XMMM's model of spontaneous mind-wandering, creative association, and daydreaming.

The threshold θ_{act} therefore defines the boundary between conscious deliberation and spontaneous cognition, and its value is modulated by the current net autonomic stress level: high stress lowers θ_{thread}, making directed goal-pursuit more persistent. Low stress (safety, satiation) raises θ_{thread}, facilitating exploratory default-mode cognition.

Threading operates on a four-mode continuum governed by $\text{Focus}(t) = \text{ES}_{\text{PD}}(t)$ (Eq. 35):

(1) *Undirected Threading* (Focus $< \theta_{\text{act}}$): all drives below threshold; Association Database traversal uses OR-gated, drive-agnostic partial similarity; the ECB is populated by recalled rather than incoming associations; corresponds to mind-wandering and sleep dreaming (motor output suppressed).

(2) *Partially Directed Threading* (Focus $\geq \theta_{\text{act}}$, loose constraint): prime drive first detected; candidate associations must share the Prime Drive ID but no further constraint; serves the goal-identification (WHAT?) function.

(3) *Directed Threading* (CS < 0.90, active PD): no high-confidence anchor-state match available; TC populated with partially matching associations under CS–AS modulation; serves the route-inference (HOW?) function.

(4) *No Threading* (CS ≥ 0.90): exact or near-exact anchor-state match found; motor commands executed directly from the Association Database without TC involvement; habit-like automatic execution.

Transition between modes is continuous, driven by rising $\text{Focus}(t)$ and falling $\text{CS}_{\text{eff}}(t)$. Biological correspondences for all four modes are documented in Document S2, Section 2.6.

Motion Algorithm: Effector Output

The Motion Algorithm translates the action candidate selected by the Linking Algorithm into physical or motor effector commands. It accepts four input types, listed in priority order:

1. *Reflex input*: reflexes triggered in M4 bypass deliberation entirely and are executed immediately at the motor level.
2. *Phobia trigger*: pre-programmed (hard-coded) avoidance responses triggered by recognised phobia stimuli.
3. *Learned Association action*: the action $\text{EM}_{\text{candidate}}(t)$ returned by Eq. 16, representing the highest-salience learned response to the current context.
4. *Tutor override*: an external tutor-issued motor command, used during the supervised learning phase. Tutor-override commands are subject to the normal satiation-event detection pathway (M8); any resulting satiation event triggers the standard IF update (M6.3), allowing the agent to learn the value of the imposed effector motions through demonstration.

The output of the Motion Algorithm is a set of actuator commands (wheel motor voltages, joint angles, speech synthesis parameters, screen output, etc.) for the current cycle.

Autonomic Stress Emotion: Mathematical Treatment

The autonomic stress drive occupies a structurally unique position in the XMMM because it is the only drive whose error signal is raised not only by direct physiological signals but also by all other drives through coupling, and by anticipatory memory recall. This makes AS the universal motivational amplifier of the architecture.

Drive-Coupled Autonomic Stress Signal

The net error signal of the autonomic stress drive at cycle t integrates contributions from all N homeostatic and allostatic drives plus memory-recalled contributions:

$$\mathrm{ES}_{\mathrm{AS}}(t) = \mathrm{clip}_{(,0,)}\,1\left\{\sum_i [w_i \cdot |D_i(t)|] + \sum_{j \in \mathcal{A}_{\mathrm{match}}} [\mathrm{IF}_j \cdot \mathrm{AS}_j]\right\} \tag{29}$$

where the first sum runs over all N drives with coupling weights w_i ($\sum w_i = 1$), and the second sum accumulates autonomic stress stored in contextually recalled associations, scaled by their Impact Factors. The $\mathrm{clip}_{(,0,)}\,1$ operator ensures $\mathrm{ES}_{\mathrm{AS}} \in [0, 1]$.

Autonomic Stress Emotion Representations

$\mathrm{DE}_{\mathrm{AS}}$ and $\mathrm{SE}_{\mathrm{AS}}$ are generated from $\mathrm{ES}_{\mathrm{AS}}$ using the same formulae as the homeostatic drives (Eqs. 6–7). They are placed in the Body Map at a diffuse visceral location (generalised thoracic/epigastric body region), consistent with the biological substrate of sympathetic arousal felt as chest tightness, racing heart, and abdominal tension.

Secondary Reinforcement via Autonomic Stress

The crucial functional consequence of the AS coupling is that any environmental cue X that is consistently observed in the Association Database during a state of elevated AS (caused by any primary drive) will itself develop the property of generating AS reduction when observed in the future — provided it has been followed, on enough occasions, by satiation. This is the mechanism of secondary reinforcement:

$$\mathrm{IF}_j \to 1 \iff A_j \text{ consistently precedes satiation events} \tag{30}$$

Once IF_j is high, mere recognition of the cue stored in A_j generates an anticipated autonomic stress reduction — a predicted $\mathrm{SE}_{\mathrm{AS}}$ — which is placed in the Body Map as a positive emotion. This is why Xzistor agents can be strongly motivated by a green door that precedes a food source, even though the green door itself provides no direct homeostatic value: the door has become a secondary reinforcer through accumulated IF weight and the consequent anticipatory AS reduction it generates. This mechanism provides a full computational account of secondary reinforcement, conditioned incentive salience, and the broad class of approach behaviours toward reward-predictive cues.

Cycle Timing and Real-Time Constraints

The complete logic loop (steps 1–12) completes within approximately 8 ms at 10 Hz for $N = 20$ drives and $|\mathcal{A}| \leq 10{,}000$ associations on a standard single-core 2.4 GHz processor. The dominant computational cost is the Association Database search ($O(|\mathcal{A}|)$ per cycle), which is reduced to $O(\log |\mathcal{A}|)$ with a B-tree index on the anchor state key. All other components (M2–M4, M8–M9) execute in $O(N)$ time.

Summary of Key Equations

Eq.	Expression	Mechanism	Section
(1)	$\mathrm{ES}_i(t) = \lvert \mathrm{CV}_i(t) - \mathrm{SP}_i \rvert \;/\; \mathrm{CV}_{\max,i}$	Error signal	M2.1
(2)	$\mathrm{ES}_{\mathrm{thirst}} = w_O O(t) + w_V V(t) + w_{\mathrm{AII}} A_{\mathrm{II}}(t)$	Multi-CV drive	M2.1
(3)	$\mathrm{ES}_{\mathrm{AS}}(t) = \max_i[w_i \mathrm{ES}_i] + \sum_j [\mathrm{IF}_j \mathrm{AS}_{\mathrm{recalled},j}]$	Autonomic stress	M2.3
(4)	$\mathrm{ES}_{\mathrm{anger}} = \min(2\mathrm{ES}_{\mathrm{AS}},\, 1) \cdot \mathbf{1}[\mathrm{ES}_{\mathrm{AS}} < 0.5]$	Anger drive	M2.3
(5)	$\mathrm{ES}_{\mathrm{fear}} = \min(2(\mathrm{ES}_{\mathrm{AS}} - 0.5),\, 1) \cdot \mathbf{1}[\mathrm{ES}_{\mathrm{AS}} \geq 0.5]$	Fear drive	M2.3
(6)	$\mathrm{DE}_i(t) = -\mathrm{ES}_i(t) \in [-1, 0]$	Deprivation emotion	M3.1
(7)	$\mathrm{SE}_i(t) = \min(\lvert \mathrm{dES}_i/\mathrm{d}t \rvert \;/\; \dot{E}S_{\max,i},\, 1) \in [0, +1]$	Satiation emotion	M3.2
(8)	$D_i(t) = \mathrm{DE}_i$ if $\mathrm{dES}_i/\mathrm{d}t \geq 0$; SE_i if $\mathrm{dES}_i/\mathrm{d}t < 0$	Signed drive value	M3.3
(9–10)	$\Omega = (r_D \times r_S)^N = (100 \times 100)^{20} = 10^{42}$	Emotion space	M3.5
(14)	$\mathrm{IF}_j(t+1) = \alpha \cdot \mathrm{IF}_j(t) + (1-\alpha) \cdot \mathrm{SE}_{\mathrm{PD}}(t)$	Reinforcement update	M5.2
(16)	$\mathrm{EM}_{\mathrm{cand}} = \arg\max_{A_j}[\mathrm{IF}_j \cdot \mathrm{sim}(\mathrm{AS}_j, \mathrm{AS}(t))]$	Action retrieval	M5.3
(18)	$\Delta\mathrm{IF}_j = \gamma^{(t^* - t_j)} \cdot \mathrm{SE}_{\mathrm{PD}}(t^*)$	Reward backprop	M6.2
(21)	$\delta(t) = \Delta\mathrm{AS}_{\mathrm{obs}}(t) - \Delta\mathrm{AS}_{\mathrm{pred}}(t)$	Prediction error	M7.1
(22)	$\mathrm{BSOR}(t) = \min(\lvert\delta\rvert/\delta_{\max},\, 1) \cdot \mathbf{1}[\lvert\delta\rvert \geq \theta_{\mathrm{BSOR}}]$	BSOR activation	M7.1
(26)	$\mathrm{PD}(t) = \arg\max_i \lvert D_i(t) \rvert$	Prime drive	M8.1
(27)	mode(t): Directed if $\lvert D_{\mathrm{PD}} \rvert \geq \theta_{\mathrm{act}}$; Default otherwise	Threading	M8.3
(29)	$\mathrm{ES}_{\mathrm{AS}}(t) = \mathrm{clip}(_,0,^) \, 1\{\sum_i [w_i \lvert D_i \rvert] + \sum_j [\mathrm{IF}_j \mathrm{AS}_j]\}$	Coupled AS	M10.1

All parameters cited in this section (α, β, γ, η, κ_+, κ_-, θ_{thread}, θ_{BSOR}, K, N, r_D, r_S) are defined at first use and listed in full in Supplementary Table S1 with their default values, ranges, and biological rationale for each choice.

Two additional equations formalise the Threading mode transition and the Focus parameter,

referenced throughout this document as S0 Eqs. 35–36:

$$\text{Focus}(t) = \text{ES}_{\text{PD}}(t) = |\text{DE}_{\text{PD}}(t)| \in [0, 1] \tag{35}$$

$$\text{mode}(t) = \begin{cases} \text{Directed} & \text{if Focus}(t) \geq \theta_{\text{act}} \\ \text{Undirected} & \text{otherwise} \end{cases} \tag{36}$$

where $\theta_{\text{act}} = 0.10$ in reference implementations, and during the satiation phase $\text{Focus}(t) = \text{SE}_{\text{PD}}(t)$. The ECB holds either real-time incoming state or a candidate Association re-evoked from the TC. The TC has maximum depth K_{TC} ($K_{\text{TC}} = 10$ for simple agents; upper bound 20 for complex implementations), distinct from the retrospective Temporal Buffer $K_{\text{TB}} = 3$–5 cycles used by RBP (§6.1).

Code availability: a reference implementation of the complete logic loop in Java is available at the Xzistor LAB Zenodo repository (`https://zenodo.org/xzistor`). The implementation used to generate the results in this paper corresponds to commit hash [to be provided on acceptance].

Document S1 — Experimental Setup, Tests and Results

The Xzistor Mathematical Model of Mind

Empirical Implementation, Test Protocols and Observed Results

Document S1 · Van Schalkwyk, R., Cook, D., Alvarez, C.E., Dehbozorgi, A. · Xzistor LAB · 2026

Complete account of both proof-of-concept implementations and the five dedicated test series. Cross-references to the formal specification use S0 prefix; biological validation uses S2.

Introduction to Experiments and Tests

The XMMM was developed as a functional principal model derived entirely from control-theoretic first principles, not from neuroanatomy, with the explicit aim of reproducing observable autonomous agent behaviours from a minimal initial endowment. Experimental validation pursues a correspondingly specific goal: to demonstrate that the hypothesised architectural mechanisms produce the predicted behaviours when implemented in physical or simulated agents, with no additional modules or specialised programming beyond the five algorithms and their interconnections.

Many tests and experiments were conducted throughout the development of Simmy and Troopy. The five dedicated test series in Section 5 represent the most systematically structured, each validating a distinct architectural mechanism. Additional emergent behaviours — daydreaming, entertainment seeking, route imagination, and combinatorially rich emotional expression — are noted in context.

Simmy (SIMAI-X1) — Virtual Agent: Experimental Setup

Simmy is a 3D virtual agent implemented in C++ and OpenGL, running the XMMM logic loop at 10 Hz. The environment is rendered in real time and the agent's face displays dynamic expressions computed from its emotion vector. Active drives include hunger, thirst, cold, warm, pain (two body locations), stress, nausea, fatigue, acute fear and anger ($N = 11$). The learning environment is a miniature virtual room containing furniture and reward and aversive prop objects. All six core mechanisms (S0 M2–M7) were validated as emergent behaviours under tutor-supervised training.

Troopy — Physical Robot: Experimental Setup

Troopy is a mobile robot based on the Lego Mindstorms platform with LeJOS open-source firmware supporting a full Java implementation of the XMMM (>30,000 lines of commented code) running at 1–3 Hz. Active drives include battery level, thirst, pain (two body locations), cold, warm, nausea, anger, acute fear, fatigue, bonding and autonomic stress ($N = 12$).

A key methodological note: the hunger drive was implemented as a time-based proxy homeostat — a control variable $\mathrm{CV}_{\mathrm{hunger}}$ rising monotonically to its maximum within approximately four

minutes, independently of any physical sensor. The executive received this as a body-located deprivation emotion DE_{hunger} indistinguishable in format from any physiologically grounded drive, demonstrating the Epistemic Isolation Principle (S0 Section 3.4): the executive acts on what it *feels*, not on the nature of the underlying control variable.

Overview of Tests Conducted

Five dedicated test series were conducted, each targeting a specific mechanism of the XMMM architecture. The tests are ordered by increasing cognitive complexity.

Dedicated Tests

Learning — Reward-based Backpropagation and Navigation Test

Introduction and Test Objectives

This test series provided the primary validation of Reward-based Backpropagation (RBP, S0 Sections 6.2–6.3). The objective was to demonstrate that an agent with no prior knowledge, no externally specified reward function, and no pre-programmed navigation behaviour could acquire goal-directed navigation through homeostatic drive-satiation reinforcement alone. A secondary objective was to demonstrate that conditioned fear responses, facial expression, daydreaming and route imagination would emerge from the same architecture without additional programming.

Test Results: Observations and Measurements

Both Simmy and Troopy were initialised with empty Association Databases. During early exploration or under virtual tutor guidance, agents encountered objects that altered their homeostatic drives. Interaction with a food reward source generated a satiation emotion SE_i that reinforced the current experience association through an increase in its Impact Factor IF_j. Tutor assistance time decreased monotonically and paths became progressively more direct.

The following additional emergent behaviours were observed without modification to the architecture: (i) conditioned fear responses to pain-associated objects, arising from allostatic AS drive re-evocation on recognition; (ii) stress-proportional facial expressions driven directly by Body Map DE_i/SE_i magnitudes; (iii) daydreaming — spontaneous undirected threading during post-satiation rest periods; and (iv) route imagination — directed threading populating the Threading Cache (TC) with high-IF landmark associations before committing to physical movement.

Conclusion and Validation Statement

The test confirms that RBP produces goal-directed learning from homeostatic reinforcement without external reward specification. The emergence of fear conditioning, graded facial expression, daydreaming and route imagination from the same architecture validates the claim that the XMMM's five-algorithm recursive loop generates a qualitatively richer behavioural repertoire than was explicitly specified.

Problem-Solving Test

Introduction and Test Objectives

This test series validated the threading mechanism (S0 Sections 8.3, 5.5) as a problem-solving substrate: the capacity of drive-motivated directed associative search to identify an appropriate action in a novel situation for which no precise Anchor State match exists in the Association Database.

Test Results: Observations and Measurements

Troopy was trained to obtain food by activating a specific-coloured button on a four-button panel; the reward-producing button was then changed without additional training. When the robot initially repeated the now-ineffective action, rising hunger deprivation $|\mathrm{DE}_{\mathrm{hunger}}|$ crossed the directed-threading threshold θ_{act}. The architecture initiated systematic associative search, retrieving the highest-IF associations relevant to the current drive state and applying their stored motor commands as candidate hypotheses. The robot subsequently identified the new reward-producing button and achieved satiation.

Conclusion and Validation Statement

The test confirms that directed threading enables inductive inference in novel sensory configurations without any explicit planning algorithm. The substrate is architecturally identical whether the novel problem is a changed button or an abstract domain such as engineering design.

Pain and Fear Test — Including Phobia Seeding

Introduction and Test Objectives

This test series validated two aspects of the XMMM's fear architecture: (i) naturally acquired fear conditioning through operant learning; and (ii) the phobia seeding procedure — installation of an innate fear response to an arbitrary sensory stimulus that has never been associated with harm. A video recording is available at `https://youtu.be/HZYeBJ6G_AE`.

Test Results: Observations and Measurements

Naturally acquired fear conditioning. After Troopy had encountered a pain-inducing floor object on multiple occasions, the object subsequently elicited avoidance behaviour, elevated AS emotion, and facial distress expression even when the pain source was not currently active. This arose from the allostatic AS drive architecture (S0 Section 2.3): stored deprivation emotions $\mathrm{DE}_{\mathrm{pain}}$ tagged to the object's Anchor State were re-evoked on recognition, without any current nociceptive signal.

Phobia seeding. A phobia of a specific novel visual stimulus — a colour never associated with any aversive event — was installed using a two-stage seeding procedure: (1) affective template acquisition from genuine noxious exposure; (2) sensory substitution replacing the pain object's Anchor State sensory component with the phobia stimulus, written to the non-modifiable database partition. Upon first encounter the agent displayed immediate intense fear responses equivalent

to those produced by direct pain. The response was not extinguishable through repeated safe encounters.

Conclusion and Validation Statement

The test confirms that the allostatic AS drive produces naturalistic conditioned fear through standard RBP, and that phobia seeding installs innate aversive responses through affective template transplantation. Both outcomes arise from the standard Association Algorithm without additional modules. The phobia seeding result provides a principled approach to value alignment in physically embodied AI agents.

Social Bonding Test

Introduction and Test Objectives

This test series validated the Social Bonding (SB) drive as a dedicated homeostatic control variable capable of generating attachment-like behaviour and emotional states through a *primary* motivational drive, without secondary reinforcement and without any modification to the core five-algorithm XMMM architecture.

The central experimental question was whether social bonding behaviour emerges from a drive whose reward source is defined exclusively by the presence of a recognisable social stimulus — independent of any co-occurrence with hunger, thirst, pain-relief, or other drive-satiation events. This design explicitly refutes the *secondary reinforcement* (*cupboard love*) account of social attachment, in which affiliation with a caregiver arises because the caregiver has been associated with primary biological rewards. In the Xzistor LAB implementation the bonding stimulus provided no food, water, or any other drive-relevant reinforcement to the robot.

A further objective was to map the computational behaviours generated by the SB homeostatic loop onto the predictions of the major empirical attachment theories (Lorenz, Bowlby, Harlow, Ainsworth) and to document which correspondences emerge purely from the architecture and which are parametric consequences of the same underlying drive equations.

Experimental Platform and Bonding Stimulus

Platform. The experimental platform was Troopy, the physical Xzistor robot described in Section 3 of this document.

Bonding stimulus. The bonding stimulus was a high-contrast blue colour panel mounted on a Mobile Feeder frame placed within Troopy's circular learning confine. The blue panel was identifiable in Troopy's camera feed as a 5-integer colour-panel Anchor State and was detected by the standard XMMM Correlation Sensitivity criterion:

$$\text{Bonding stimulus identified:} \quad \mathrm{CS}_{\mathrm{SB}}(t) \;\geq\; \theta_{\mathrm{CS}} = 0.90 \tag{S1.1}$$

Physical proximity to the bonding source was confirmed by a downward-facing colour sensor on Troopy's chassis detecting a green floor panel positioned beneath the Mobile Feeder (the *contact*

zone), providing a binary contact signal analogous to body-contact comfort. A black floor band immediately beyond the green zone triggered the pain homeostatic drive at maximum intensity ($\mathrm{ES}_{\mathrm{pain}} = 1.0$), used in the pain-dominance sub-test (Section 5.4.5).

The Mobile Feeder provided *no food, water, or any other drive reinforcement* in the baseline social bonding tests. This is the operational implementation of the primary-motivation hypothesis.

Social Bonding Drive Specification

The SB drive is implemented as a time-dependent homeostatic accumulation-and-decay control loop. During deprivation ($\mathrm{CS}_{\mathrm{SB}}(t) < \theta_{\mathrm{CS}}$), the error signal increases linearly with elapsed absence time:

$$\mathrm{ES}_{\mathrm{SB}}(t) \;=\; \min\!\left(1,\; \kappa^{\mathrm{SB}}_{\mathrm{dep}} \cdot t_{\mathrm{abs}}(t)\right) \tag{S1.2}$$

where $\mathrm{ES}_{\mathrm{SB}}(t) \in [0, 1]$ is the normalised SB drive error signal, $\kappa^{\mathrm{SB}}_{\mathrm{dep}} > 0$ is the deprivation accumulation rate constant (normalised drive units per second), and $t_{\mathrm{abs}}(t)$ is the elapsed time in seconds since $\mathrm{CS}_{\mathrm{SB}}$ last satisfied $\geq \theta_{\mathrm{CS}}$.

When the bonding stimulus is positively identified *and* the contact zone is active, the drive enters satiation and decays exponentially:

$$\frac{d\mathrm{ES}_{\mathrm{SB}}}{dt}(t) \;=\; -\kappa^{\mathrm{SB}}_{\mathrm{sat}} \cdot \mathrm{ES}_{\mathrm{SB}}(t), \quad \mathrm{CS}_{\mathrm{SB}}(t) \geq \theta_{\mathrm{CS}} \text{ and floor contact} \tag{S1.3}$$

The corresponding somatosensory Body Map representations are:

$$\mathrm{DE}_{\mathrm{SB}}(t) = -\mathrm{ES}_{\mathrm{SB}}(t) \;\in [-1, 0] \qquad \text{(separation distress; aversive, body-located)} \tag{S1.4}$$

$$\mathrm{SE}_{\mathrm{SB}}(t) = 1 - \mathrm{ES}_{\mathrm{SB}}(t) \;\in [0, 1] \qquad \text{(connection; appetitive, body-located)} \tag{S1.5}$$

Autonomic Stress coupling follows S0 Eq. 3:

$$\Delta\mathrm{ES}_{\mathrm{AS}}(t) \;=\; \alpha_{\mathrm{SB}} \cdot \mathrm{ES}_{\mathrm{SB}}(t) \tag{S1.6}$$

producing the generalised anxious arousal of social separation. The distress vocalisation (learn-modifiable crying reflex) fires when:

$$\mathrm{ES}_{\mathrm{SB}}(t) \;>\; \theta_{\mathrm{cry}} \tag{S1.7}$$

The SB drive becomes Prime Drive when $\mathrm{ES}_{\mathrm{SB}}(t)$ is the largest normalised error signal across all active drives and $\mathrm{ES}_{\mathrm{SB}}(t) \geq \theta_{\mathrm{act}} = 0.10$. No modification to any other component of the five-algorithm architecture was required. A full biological correspondence mapping is provided in Document S2, Section 1.4.

Experimental Configurations

Two test configurations were implemented sequentially, differing in the relative dominance and accumulation rate of the SB drive and in the presence or absence of innate navigational reflexes.

Parameter	Test Case 1: High Dominance	Test Case 2: Naturalistic
$\kappa_{\text{dep}}^{\text{SB}}$	Very high: ES_{SB} reaches 1.0 within 5 s of absence	Moderate: ES_{SB} reaches 0.30 over 3 min
Max ES_{SB}	1.0	0.30
θ_{act}	0.10	0.10
θ_{cry}	0.30	0.20
Innate navigation reflex	Active: rotational search until $\text{CS}_{\text{SB}} \geq 0.90$, then approach	Absent: distress signal only; navigation requires tutor
Contact requirement	Green floor panel + blue panel recognition	Green floor panel + blue panel recognition
SB vs. other drives	Dominant: SB > all other drives at all times	Subordinate: hunger, thirst, pain and fear take precedence
Mobile Feeder as food	No (baseline); yes in compound sub-test	No (baseline); yes in compound sub-test
Pain sub-test	Black band beyond green zone; pain reflex active	Tested
Runs to autonomy	≈50 tutor-assisted runs	≈50 tutor-assisted runs (slower; drive competition)

Test Results: Test Case 1 (High-Dominance Configuration)

Post-initialisation behaviour. Upon system initialisation ('birth'), $\text{ES}_{\text{SB}}(t)$ reached 1.0 within five seconds, immediately making the SB drive the Prime Drive. The innate rotational search reflex activated without prior learning: Troopy rotated until $\text{CS}_{\text{SB}}(t) \geq 0.90$, then executed the approach motor programme. Upon contact (green floor panel detected), the satiation event fired and ES_{SB} and the coupled ES_{AS} declined together. This replicates neonatal rooting and approach behaviour, which is similarly innate and precedes learned navigation in mammalian infants.

Pain sub-test: drive dominance under aversive stimulation. The involuntary stop-and-contact reflex was suppressed and a pain-triggering black floor band was placed immediately beyond the green floor panel. Each approach triggered maximum-intensity pain activation ($\text{ES}_{\text{pain}} = 1.0$) and the involuntary retreat reflex. Despite this, Troopy consistently re-approached the Mobile Feeder: the dominant SB drive maintained proximity-seeking in the presence of repeated aversive consequences. The fear conditioning that normally inhibits approach to pain-associated objects (as documented in Section 5.3) was overridden by SB drive dominance. This result demonstrates that:

1. The SB drive generates approach behaviour toward its reward source independently of the reward source's other drive-relevant properties.

2. Drive dominance parameters directly control the agent's approach–avoidance balance near a mixed-valence attachment figure.

3. A sufficiently dominant SB drive constitutes an architectural safety property: the agent maintains proximity to its bonding object even at personal cost, without external instruction.

Existential drive conflict. In the pure high-dominance configuration without compound reward, the SB drive recaptured Prime Drive status within five seconds of any disengagement, preventing the hunger drive from sustaining Prime Drive status long enough to complete navigational training trials. This blocked acquisition of autonomous hunger navigation competency — an 'existential conflict' with a direct biological analogue: a neonate that cannot tolerate brief separation from its caregiver cannot develop independent foraging competency. This observation motivated the naturalistic configuration (Test Case 2).

Compound reward sub-test. When the Mobile Feeder was additionally configured as a hunger reward source, Troopy exhibited sustained proximity maintenance, simultaneously satisfying the SB drive (visual recognition) and the hunger drive (physical contact), replicating the functional relationship between infant and lactating mother. Removal of the innate navigation reflex under these conditions produced distress vocalisation followed by quiescence upon tutor repositioning of the Mobile Feeder — reproducing the neonatal pattern in which the infant signals distress to elicit caregiver approach rather than independently navigating.

Test Results: Test Case 2 (Naturalistic Configuration)

Drive hierarchy and survival-critical prioritisation. With $\kappa_{\text{dep}}^{\text{SB}}$ reduced to produce a maximum ES_{SB} of 0.30 over three minutes, the SB drive became Prime Drive only after hunger, thirst, pain, and fear drives had been satiated. Troopy navigated to the food source when hungry, to the water source when thirsty, and to the Mobile Feeder episodically when the SB drive reached Prime Drive status. After approximately fifty tutor-assisted training runs, Troopy achieved fully autonomous SB-driven navigation, consistent with the learning trajectory observed in the navigation and problem-solving tests. The reduced SB dominance also resolved the existential conflict observed in Test Case 1: the agent could now sustain hunger-navigation training sessions between episodes of social bonding drive satiation.

Impact Factor modulation by drive context. When the Mobile Feeder was configured as a combined SB and hunger reward source, the compounded satiation events across both drives produced systematically higher Impact Factor values in the Associations formed during Mobile Feeder proximity. This formalises the well-documented observation that primary caregivers who simultaneously provide nourishment and physical comfort receive more stable and intense attachment than those providing comfort alone. The converse was equally demonstrated: repeated pain conditioning from the black floor band progressively reduced the Mobile Feeder's net Impact Factor, and under sufficiently severe conditioning the fear/pain drive exceeded the SB drive, rendering the bonding object a net aversive stimulus. The agent avoided it — a computational

replication of fear-conditioned attachment inhibition as documented in Harlow's (1958) rhesus experiments.

Safe-base exploration: emergent attachment dynamic. The most behaviourally compelling emergent phenomenon was spontaneous approach–avoidance oscillation. Having satiated the SB drive, Troopy initiated exploratory trajectories into the confine. As distance from the Mobile Feeder increased, $\mathrm{ES}_{\mathrm{SB}}(t)$ and the coupled $\mathrm{ES}_{\mathrm{AS}}(t)$ rose. Upon detecting Prime Drive pressure, Troopy navigated back to the bonding source, experiencing a compound satiation event — dual relief of SB deprivation and accumulated AS — upon recognition and contact.

This oscillation was not programmed. It emerged from three architectural elements: (1) time-dependent SB deprivation accumulation (Eq. S1.2); (2) coupled AS rise during absence; and (3) RBP-reinforced navigational competency routing the agent back when SB was Prime Drive. Critically, Troopy learned that departing and returning generated a larger total satiation event than remaining stationary in proximity, because satiation magnitude is proportional to $\mathrm{ES}_{\mathrm{SB}}$ at the moment of recognition. The agent therefore initiated departures with *progressively increasing duration and distance* — the computational realisation of Bowlby's (1969) secure-base phenomenon: the attached agent uses the bonding source as a base from which to explore, periodically returning for re-motivation, with the exploration radius self-organising via the reward structure of delayed reunion.

Qualitative behavioural sequence (observed).

1. SB drive crosses $\theta_{\mathrm{act}} = 0.10$; AS rises in concert.
2. Distress vocalisation activated ($\mathrm{ES}_{\mathrm{SB}} > 0.20$); tutor navigates Troopy to Mobile Feeder (initial training runs).
3. $\mathrm{CS}_{\mathrm{SB}} \geq 0.90$ and floor contact: satiation event. $\mathrm{ES}_{\mathrm{SB}}$ and $\mathrm{ES}_{\mathrm{AS}}$ decline; BSOR_{+} fires; preceding Associations receive Impact Factor reinforcement.
4. Post-satiation: all drives below θ_{act}; Troopy enters undirected threading (exploratory play / daydreaming).
5. As t_{abs} increases, $\mathrm{ES}_{\mathrm{SB}}$ and $\mathrm{ES}_{\mathrm{AS}}$ rise; Threading Cache populates with SB-satiation-tagged Mobile Feeder Associations; autonomous return navigation begins.
6. Reunion satiation magnitude proportional to accumulated $\mathrm{ES}_{\mathrm{SB}}$; longer departures produce larger satiation events, reinforcing progressively longer exploratory sorties.

Conclusion and Validation Statement

The social bonding test series validates four architectural claims:

1. *Primary motivation.* Attachment-like proximity-seeking and contact comfort emerge from a dedicated homeostatic drive without secondary reinforcement from any other drive, directly refuting the cupboard-love account within the XMMM framework.

2. *Emergent ethological correspondences.* All four classes of attachment behaviour — proximity-seeking, contact comfort, pain-tolerant approach, and safe-base exploration — are direct consequences of the SB drive equations and the XMMM's standard RBP mechanism; none was explicitly programmed.

3. *Drive-parametric attachment styles.* The two test configurations produce behaviourally distinct profiles corresponding to the high-dependency neonatal pattern (Test Case 1) and the secure-base adult pattern (Test Case 2), demonstrating that the full developmental trajectory of attachment is parameterised within the single SB drive architecture.

4. *Architectural safety property.* The pain-dominance sub-test demonstrates that an SB drive calibrated above all other drives constitutes an architecturally guaranteed proximity-preservation property toward the bonding object — an approach to AI alignment by motivational necessity rather than externally imposed constraint.

Biological correspondence analysis, including the complete mapping of SB drive equations onto the eight-component hypothalamic-amygdala-VTA circuit, is provided in Document S2, Section 1.4. Post-hoc correspondence with the attachment theories of Lorenz, Bowlby, Harlow, and Ainsworth is provided in Table S2.1 of Document S2.

Animal versus Human Brain — Threading Dissociation Test

Introduction and Test Objectives

This test series constitutes a parametric lesion study: the threading mechanism of the Linking Algorithm was programmatically disabled in Troopy while all other components — the Drive, Reflex and Association Algorithms and RBP — remained fully active. The objective was to isolate the contribution of directed threading from operant learning within the same agent, same trained Association Database, and identical environmental conditions.

Test Results: Observations and Measurements

With threading disabled, the robot continued to learn through RBP and, when presented with a previously trained sensory context at high fidelity, immediately retrieved and executed the correct motor sequence. Associative recognition-execution was fully intact. The deficit appeared exclusively when the environment had changed in a way that placed the agent in a partially familiar but not directly learned configuration. In this condition the threading-disabled robot showed complete cessation of goal-directed navigation, became progressively more deprived (rising $|\mathrm{DE}_{\mathrm{hunger}}|$), and defaulted to a pre-programmed hunger distress reflex.

Conclusion and Validation Statement

The test establishes a clean dissociation: associative recognition-execution (intact without threading) versus directed inductive inference (absent without threading). This behavioural profile parallels the reactive, cue-dependent behaviour characterising non-human animals in novel or ambiguous environments. Full biological correspondence analysis is provided in Document S2, Section 2.

Experimentation: Future Directions

The five test series above provide proof-of-concept validation of the XMMM's core mechanisms under naturalistic conditions. The following controlled validation experiments are identified as the highest-priority next research directions.

Controlled Navigation Study

Systematic replication of the navigation test with matched start positions, recording time-to-reward and approach path length across a defined number of trials, to characterise the IF accumulation rate and its dependence on RBP parameters (S0 Eqs. 14–15).

Social Bonding Control Experiment (Pre-registered)

Two agent cohorts trained with reversed bonding-stimulus assignments: a blue colour panel serves as the bonding stimulus for cohort A; a green colour panel serves as the bonding stimulus for cohort B. Both panels are physically present throughout all trials for all agents. Exclusive preference for the assigned stimulus in bonding-deprivation probe trials, irrespective of which panel is *not* the trained bonding stimulus, will confirm that the preference arises from the SB drive's homeostatic reinforcement history rather than from any stimulus-intrinsic property. This experiment also controls for any differential salience of the two colour panels.

Phobia Extinction Resistance

Systematic test of the non-modifiable partition property through repeated safe exposure to the phobia stimulus, with the number of safe trials required for measurable IF decay recorded as a quantitative outcome.

Deliberative Capacity Scaling

A systematic Focus clamping study across Focus values 0.0 to 1.0 in steps of 0.1, using the same agent, Association Database and environmental configuration. This tests the predicted linear relationship between ES_{PD} and goal-directed performance specified in S0 Eq. 35.

Language Acquisition

An initiated project in which an XMMM agent interacts with a human tutor whose provision of satiation is contingent on specific communicative acts, testing whether consistent communicative patterns emerge from drive-motivated reinforcement without a dedicated language faculty.

Document S2 — Biological Validation and Correspondences

The Xzistor Mathematical Model of Mind

Empirical Grounding, Drive Validation and Neuroanatomical Correspondence

Document S2 · Van Schalkwyk, R., Cook, D., Alvarez, C.E., Dehbozorgi, A. · Xzistor LAB · 2026

Section 1 provides biological validation of XMMM drive homeostats. Section 2 provides post-hoc neuroanatomical correspondence. Reference numbering continues from the main manuscript (refs 1–35); new references are numbered 36 onwards.

Biological Validation of XMMM Drive Homeostats

The XMMM's drive homeostats are numerical control loops: simplified formal models of biological drive systems that capture the information-processing relationships between sensory input, set-point comparison, drive error generation, and executive output, without modelling the full biochemical, cellular, or circuit-level complexity of their biological counterparts. The validations in this section demonstrate that each simplified homeostat nevertheless produces behaviours that are consistent with those produced by the biological system — and, where tested, that the functional constraints required for biologically plausible behaviour in the XMMM converge on independently established biological facts. This convergence is the appropriate test of a functional principal model: not neuroanatomical fidelity, but functional equivalence at the level of information processing that governs observable behaviour.

Thirst Drive Validation

The XMMM thirst homeostat was the first to be systematically mapped to its biological counterpart by Van Schalkwyk and Cook (2024). The multi-variable biological thirst drive — integrating plasma osmolality $O(t)$, blood volume $V(t)$, and angiotensin II concentration $A_{\text{II}}(t)$ — is modelled in the XMMM as a weighted linear combination (S0 Eq. 2):

$$\text{ES}_{\text{thirst}}(t) = w_O \cdot O(t) + w_V \cdot V(t) + w_{\text{AII}} \cdot A_{\text{II}}(t) \tag{S2.1}$$

The mapping proceeds as follows. The subfornical organ (SFO) and organum vasculosum laminae terminalis (OVLT) serve as the primary osmolality and volume sensors, implementing the $\text{CV}(t)$ measurement stage of the XMMM drive equation. The median preoptic nucleus (MnPO) integrates signals from both structures and compares them against the homeostatic set-point, implementing the $\text{ES}(t)$ computation. Projection neurons from the MnPO to the insular cortex carry the resulting drive error signal to the interoceptive representation stage, implementing the Body Map placement of $\text{DE}_{\text{thirst}}$. The subjective experience of thirst — a body-located aversive sensation in the oral-pharyngeal and epigastric regions — is the biological implementation of this DE placement. The hypothalamic oxytocin projection system and vasopressin-mediated peripheral water conservation implement the satiation arm, generating $\text{SE}_{\text{thirst}}$ as the rate of drive error reduction following drinking. Zimmerman et al. (2016)[22] demonstrated that hypothalamic thirst neurons anticipate the homeostatic consequences of drinking before blood osmolality

changes, consistent with the XMMM's gustatory delay implementation (see Section 1.1.1).

Gustatory Delay: Convergent Validation

During early implementation of the Troopy robot, an oscillatory behaviour was observed: brief food contact reduced the hunger control variable just below the detection threshold, causing the agent to disengage and move away, whereupon the drive immediately resumed its upward trajectory and pulled the agent back — trapping it in a repetitive approach-retreat cycle. Resolving this required the incorporation of a gustatory delay: a period during which the satiation signal persists after food contact before the control variable resumes its upward trajectory.

This architectural requirement converges independently with a well-established biological fact. In mammals, oropharyngeal and gastrointestinal metering systems maintain a satiation signal for approximately 15 minutes after food or water intake, well before blood glucose or plasma osmolality register a measurable change[22]. The XMMM's functional requirement for a temporal delay, derived from observing physically implausible oscillatory behaviour in the robot implementation, thus independently converged on a biological mechanism that was not part of the model's initial specification. The fact that the XMMM had no prior knowledge of oropharyngeal metering latency, yet required precisely such a mechanism for biologically realistic behaviour, constitutes an independent validation of the Drive Algorithm's control-theoretic structure.

Pain Drive Validation

The XMMM pain homeostat models nociceptive input as a localised body-map drive, with $\mathrm{CV}_{\mathrm{pain}}(t)$ rising upon nociceptive contact and falling upon removal of the pain stimulus. The resulting $\mathrm{DE}_{\mathrm{pain}}$ is placed at a specific Body Map location corresponding to the anatomical site of the pain event, implementing the sensory-discriminative component of pain. The AS drive coupling to $\mathrm{DE}_{\mathrm{pain}}$ implements the affective-motivational component: pain is aversive not merely because it is a drive error but because it co-activates the autonomic stress drive, generating a composite aversive state that motivates immediate corrective action.

This two-component structure — sensory-discriminative (localised body-map DE) and affective-motivational (AS coupling) — maps directly onto the biological distinction between the lateral spinothalamic tract, which carries sensory-discriminative nociceptive information to somatosensory cortex, and the medial spinothalamic tract, which carries affective-motivational pain signals to the anterior cingulate cortex, insular cortex and amygdala[16]. In the XMMM, these two pathways correspond to the direct DE placement in the Body Map (lateral pathway: location and intensity) and the AS coupling that generates the aversive motivational urgency (medial pathway: suffering and the drive to act). This post-hoc convergence — a two-component pain architecture arrived at independently from control-theoretic principles — provides further evidence that the XMMM's functional decomposition captures the essential computational organisation of the biological pain system.

Autonomic Stress Drive

The Autonomic Stress (AS) drive occupies a unique architectural position: it activates in parallel with every other homeostatic and allostatic drive, amplifying motivational urgency and generating the aversive subjective quality — the felt urgency of a need — that makes deprivation states genuinely unpleasant rather than merely informational. It is also the mechanism through which memory-recalled associations can generate full autonomic responses in the absence of any current physiological threat.

Autonomic Stress Coupling Validation

The AS coupling equation (S0 Eq. 3) produces a consolidated net stress signal from all concurrent drive deprivations plus recalled associative AS contributions. This synchronous coupling of AS with all homeostatic drives maps onto the known hypothalamic-pituitary-adrenal (HPA) axis co-activation with homeostatic drive centres: the paraventricular nucleus of the hypothalamus receives input from virtually all homeostatic monitoring structures and releases CRH, driving pituitary ACTH and adrenal cortisol release in proportion to cumulative homeostatic load[11]. Benarroch's (1993) systematic account of hypothalamic integrative functions — published after the XMMM's 2002 specification — describes precisely the multi-drive convergent activation structure that the XMMM AS coupling formalises mathematically.

Anger and Acute Fear as Derived Drives

The XMMM derives anger and acute fear from the autonomic stress error signal using a piecewise linear mapping (S0 Eqs. 4–5). Let $\mathrm{ES}_{\mathrm{AS}}(t) \in [0, 1]$ be the normalised autonomic stress signal:

$$\mathrm{ES}_{\mathrm{anger}}(t) = \min\big(2 \cdot \mathrm{ES}_{\mathrm{AS}}(t),\ 1.0\big) \cdot \mathbf{1}\big[\mathrm{ES}_{\mathrm{AS}} < 0.5\big] \tag{S2.2}$$

$$\mathrm{ES}_{\mathrm{fear}}(t) = \min\big(2 \cdot (\mathrm{ES}_{\mathrm{AS}}(t) - 0.5),\ 1.0\big) \cdot \mathbf{1}\big[\mathrm{ES}_{\mathrm{AS}} \geq 0.5\big] \tag{S2.3}$$

Anger is therefore active across the lower half of the autonomic stress range (0–50%), motivating assertive approach-and-remove behaviour toward the stressor. Acute fear activates in the upper half (50–100%), motivating withdrawal and evasion. This piecewise structure was derived empirically from observing the natural breakpoints of aggressive and avoidance behaviour in the robot implementations — in the absence, at the time of specification, of clear quantitative biological data on anger-fear transition thresholds. Post hoc, this approach-withdrawal transition is consistent with the known functional organisation of the amygdala and periaqueductal grey (PAG): the basolateral amygdala and ventral PAG support active, approach-oriented threat responses under moderate threat, while the central amygdala and dorsal PAG generate passive, withdrawal-oriented responses under extreme threat[26]. The XMMM's piecewise transition at 50% AS corresponds functionally to this BLA–CeA/vPAG–dPAG functional switch, providing a retrospective biological grounding for an engineering decision made purely on the basis of observed robot behaviour.

Social Bonding: Biological Homeostat and XMMM Correspondence

Social bonding occupies a special position in the XMMM's drive architecture. Unlike physiological drives such as hunger and thirst, whose biological counterparts are well-characterised single-organ systems, the social bonding homeostat in the brain is a multi-component circuit spanning the hypothalamus, amygdala, ventral tegmental area, nucleus accumbens, medial prefrontal cortex, hippocampus, and HPA axis. Despite this biological complexity, the information that reaches the executive controller is a single valenced signal: the felt experience of loneliness (deprivation) or belonging (satiation). The XMMM's social bonding drive formalises precisely this sub-executive data compression.

The Biological Social Bonding Homeostat

The biological social bonding system has the structure of a homeostatic control loop. Lee et al. (2021)[27] proposed that evolutionarily conserved neural systems underlie the maintenance of social connections, alert the individual to their absence, and coordinate effector mechanisms to restore social contact — a circuit organisation explicitly described as homeostatic and directly paralleling the systems regulating energy and fluid balance. Huang et al. (2025)[28] identified two previously uncharacterised neuronal populations in the hypothalamic medial preoptic nucleus that are specifically activated during social isolation and social rebound respectively, encoding social need and social satiety in a manner directly analogous to the AgRP and POMC populations governing hunger and food satiety in the arcuate nucleus. This discovery establishes that social bonding has the same anatomical specificity and mechanistic organisation as the physiological drives the XMMM models most explicitly.

The full biological circuit encompasses eight anatomically distinct components, each mapping onto a specific XMMM functional element:

Hypothalamic paraventricular and supraoptic nuclei (PVN/SON). Primary source of oxytocin, the principal satiation neuromodulator. Magnocellular PVN neurons project broadly to the amygdala, hippocampus, VTA, nucleus accumbens and prefrontal cortex[29], allowing targeted oxytocin release in circuits relevant to social bonding, reward and social memory. In the XMMM, this corresponds to the SE_{SB} satiation signal and its propagation to the Body Map and Association Algorithm.

Hypothalamic medial preoptic nucleus (social need and satiety neurons). The newly identified neuronal populations encoding social deprivation and social satiety[28] map directly onto the XMMM Drive Algorithm: the social need neurons correspond to $CV_{SB} < SP_{SB}$ (deprivation), the social satiety neurons to $CV_{SB} \geq SP_{SB}$ (satiation).

Central and medial amygdala (CeA/MeA). Serve as the primary autonomic stress coupling relay: social deprivation activates the amygdala and sympathetic nervous system, generating the aversive autonomic arousal that constitutes the subjective experience of loneliness — elevated cortisol, increased heart rate, social hypervigilance. This is the direct biological implementation of the XMMM's AS drive coupling to DE_{SB}. Oxytocin signalling at the amygdala during social reunion attenuates this stress response through the prelimbic prefrontal cortex —

implementing the AS uncoupling that the XMMM models as the transition from DE_{SB} to SE_{SB} at the satiation event[28].

Bed nucleus of the stria terminalis (BNST). Sustains the anxiety component of the autonomic stress response under chronic social deprivation, corresponding to the XMMM's allostatic AS elevation that persists during prolonged SB drive deprivation.

VTA–nucleus accumbens dopaminergic pathway. Activated by social interaction and promotes social reinforcement learning through dopamine release at the nucleus accumbens, interacting with oxytocin through the VTA–NAc pathway to produce the rewarding valence of social contact[29]. This implements Reward-based Backpropagation (RBP) at the social satiation event: the dopamine prediction error signal reinforces the associations that preceded social contact, increasing their Impact Factors IF_j.

Medial prefrontal cortex. Encodes social rank and quality assessments, implementing a modulation of the Correlation Sensitivity (CS) parameter by social context: the social quality of an interaction affects the threshold at which the bonding drive responds to a given social cue.

Hippocampal CA2 region. Specifically involved in social memory and social identity recognition[28], enabling the agent to distinguish familiar from unfamiliar conspecifics. This implements the identity component of the Anchor State in the Association Algorithm: ES_{SB} rises when a recognised familiar bonding stimulus is detected, not when any instance of the stimulus class is present.

HPA axis (cortisol). Generates sustained cortisol release under chronic social deprivation, driving the allostatic component of the autonomic stress drive and producing the well-documented health consequences of chronic loneliness.

The XMMM Social Bonding Drive: Data Compression

The XMMM reduces this eight-component biological system to a single homeostatic control variable governed by time-dependent accumulation-and-decay equations. The physical robot experiments (Document S1, Section 5.4) validated this reduction in Troopy using a high-contrast blue colour panel as the bonding stimulus — detected by the standard XMMM Correlation Sensitivity criterion ($CS_{SB} \geq \theta_{CS} = 0.90$) — and a green floor panel as the physical contact proxy. The bonding stimulus provided no food, water, or any other drive reinforcement: the SB drive was validated as a *primary* motivational system, not a secondary reinforcer.

The formal drive specification is:

$$ES_{SB}(t) \;=\; \min\!\big(1,\; \kappa^{SB}_{dep} \cdot t_{abs}(t)\big), \qquad CS_{SB}(t) < \theta_{CS} \tag{S2.4a}$$

$$\frac{dES_{SB}}{dt}(t) \;=\; -\kappa^{SB}_{sat} \cdot ES_{SB}(t), \qquad CS_{SB}(t) \geq \theta_{CS} \text{ and floor contact} \tag{S2.4b}$$

where $\kappa^{SB}_{dep} > 0$ governs the rate of social deprivation accumulation during absence of the bonding stimulus, $\kappa^{SB}_{sat} > 0$ governs the rate of satiation decay upon recognition and contact, and $t_{abs}(t)$ is the elapsed absence time in seconds. This formulation is time-dependent rather than event-count-dependent, reflecting the continuous nature of social absence as a drive signal. An

equivalent exponential approximation $\mathrm{ES}_{\mathrm{SB}}(t) = 1 - e^{-\kappa_{\mathrm{dep}}^{\mathrm{SB}} \cdot t_{\mathrm{abs}}(t)}$ converges on Eq. (S2.4a) for small $\kappa_{\mathrm{dep}}^{\mathrm{SB}} \cdot t_{\mathrm{abs}}$; the linear accumulation form was used in the physical implementation for computational simplicity.

The resulting Body Map representations are:

$$\mathrm{DE}_{\mathrm{SB}}(t) = -\mathrm{ES}_{\mathrm{SB}}(t) \ \in [-1, 0] \qquad \text{(separation distress; aversive, body-located)} \tag{S2.5}$$

$$\mathrm{SE}_{\mathrm{SB}}(t) = 1 - \mathrm{ES}_{\mathrm{SB}}(t) \ \in [0, 1] \qquad \text{(belonging; appetitive, body-located)} \tag{S2.6}$$

Autonomic Stress coupling (S0 Eq. 3) produces the physically felt aversive quality of social deprivation:

$$\Delta\mathrm{ES}_{\mathrm{AS}}(t) \ = \ \alpha_{\mathrm{SB}} \cdot \mathrm{ES}_{\mathrm{SB}}(t) \tag{S2.7}$$

The distress vocalisation (learn-modifiable crying reflex) is triggered when:

$$\mathrm{ES}_{\mathrm{SB}}(t) \ > \ \theta_{\mathrm{cry}} \tag{S2.8}$$

Reference implementation parameters validated in the Troopy experiments were: Test Case 1 (high dominance) $\kappa_{\mathrm{dep}}^{\mathrm{SB}}$ set such that $\mathrm{ES}_{\mathrm{SB}}$ reaches 1.0 within 5 s, $\theta_{\mathrm{cry}} = 0.30$; Test Case 2 (naturalistic) $\mathrm{ES}_{\mathrm{SB}}$ reaches 0.30 over 3 min, $\theta_{\mathrm{cry}} = 0.20$; both configurations: $\theta_{\mathrm{CS}} = 0.90$, $\theta_{\mathrm{act}} = 0.10$, $\alpha_{\mathrm{SB}} = 1.0$ (full AS coupling).

Data compression interpretation.. This compression is not an oversimplification but a formally correct model of the computational role the social homeostat plays at the executive interface. The Linking Algorithm does not receive oxytocin concentrations, amygdala activation magnitudes, hypothalamic preoptic neuron firing rates, or social utility scores. It receives, after sub-executive compression by the eight-component biological circuit: "I feel lonely" ($\mathrm{DE}_{\mathrm{SB}}$, rising aversive somatosensory state) or "I feel connected" ($\mathrm{SE}_{\mathrm{SB}}$, falling aversive/ appetitive state). The entire biological circuit operates sub-executively, collapsing its multi-dimensional social assessment into a single actionable felt state before it reaches the Linking Algorithm — the most concrete illustration in the XMMM of the general data compression principle described in the main manuscript, Section 2.1.

The BSOR_{+} amplification at the satiation event makes social reunion one of the most powerfully reinforcing events in the architecture, corresponding to the oxytocin/dopamine co-release documented at reunion in bonded biological dyads[29], and producing the Impact Factor reinforcement of approach associations that underlies the progressive navigational competency observed across training runs.

Empirical grounding and post-hoc attachment theory correspondences.. The equations (S2.4a)–(S2.8) were validated in the physical Xzistor robot experiments documented in Document S1, Section 5.4. All four classes of emergent behaviour — proximity-seeking, contact comfort, pain-tolerant approach under drive dominance, and safe-base exploration with progressively extending

departure distance — arose without explicit programming, as direct computational consequences of the SB drive equations interacting with the standard XMMM RBP mechanism.

Post-hoc analysis reveals that these behaviours map onto the foundational predictions of the major empirical attachment theorists. Table S2.1 provides a systematic correspondence mapping. All correspondences were identified after the experimental design was completed and were not used as design inputs.

Theorist	Key Claim	XMMM Correspondence
Lorenz	Imprinting: bonding object determined by first recognised moving stimulus in the post-hatching critical period; permanent and extinction-resistant.	SB drive reference representation set to first recognised stimulus satisfying a minimum movement criterion post-initialisation. High early IF advantage from the high-κ_{dep} neonatal phase (Test Case 1) creates permanence and extinction resistance.
Bowlby	Attachment is a primary motivational system (not secondary reinforcement). Secure base: attachment figure used as a base for exploration with periodic return for remotivation.	SB drive is a dedicated homeostatic loop independent of all other drives; validated by Mobile Feeder providing no food or water reinforcement. Safe-base exploration emerged spontaneously in Test Case 2 from the reward structure of delayed reunion satiation, without programming.
Harlow	Contact comfort preferred over feeding; drive primacy over secondary reinforcement; fear-conditioned approach inhibition can be overridden by a dominant attachment drive.	Mobile Feeder provides no food (refutes cupboard love). Pain sub-test: dominant SB drive overrides fear conditioning. Weak SB drive: pain conditioning reduces bonding IF below pain drive threshold, producing avoidance — replicating fear-conditioned attachment inhibition in Harlow's rhesus experiments.
Ainsworth	Secure: confident exploration, moderate separation distress, rapid reunion comfort. Anxious-avoidant: reduced approach, suppressed distress. Ambivalent: heightened distress, ineffective comfort.	Test Case 2 (naturalistic) produces the secure profile. Weak SB + moderate pain conditioning predicts the avoidant profile via IF reduction. Alternating pain/comfort from the bonding object predicts the ambivalent profile via concurrent SB satiation and pain/fear deprivation, with approach drive and retreat reflex competing at each contact event.

Table S2.1. Post-hoc correspondence with attachment theory. All correspondences derived after experimental design; none used as design inputs.

Safety Engineering Implication

A finding of particular relevance to AI safety emerged from the social bonding test (Document S1, Section 5.4): the pain-dominance sub-test demonstrated that an SB drive calibrated above all other drives constitutes an architecturally guaranteed proximity-preservation property toward the bonding object, even under repeated aversive experience from that object. When the bonding object is defined as a human caregiver — or, at higher abstraction, as the class "human agent" — this property produces an artificial agent that, by architectural necessity rather than externally imposed constraint, volitionally acts to maintain the presence and wellbeing of its human attachment object, because that presence constitutes the agent's most reliable homeostatic satiation source.

By calibrating the SB drive strength above the agent's physiological survival drives, a strong safety case can be constructed that such an agent will always act protectively toward humans — not because it has been instructed to do so, but because human proximity is its primary homeostatic need. This constitutes an approach to AI alignment by motivational necessity, robust to the failure modes that affect rule-based and reinforcement-learning-based alignment approaches, because it operates at the level of innate motivational architecture.

Functional Correspondences with Brain Neuroanatomy

The correspondences in this section map each XMMM functional algorithm onto its identified biological correlate. Because the XMMM was derived entirely from control-theoretic first principles without neuroanatomical data as input, each correspondence documented below is a post-hoc convergence: an independent derivation that arrived at the same functional organisation as the biological brain. The fact that this convergence holds across all six algorithms and the Linking Algorithm — without exception — constitutes the strongest class of theoretical validation available to a functional cognitive architecture.

Body Map — Insular Cortex

The XMMM Body Map places body-located, magnitude-coded, valence-signed emotional representations in a data structure that the Linking Algorithm reads as its primary motivational input. The biological structure implementing this function is the insular cortex.

Craig (2009)[14] established that the posterior-to-anterior gradient of the insular cortex performs a hierarchical re-representation of homeostatic signals: raw visceral inputs arrive at the posterior insula and are progressively integrated, re-coded, and assigned motivational valence as information flows anteriorly, culminating in subjective feeling states at the anterior insula. This is functionally identical to the XMMM's Drive Algorithm → Body Map transform: $\mathrm{CV}(t) \rightarrow \mathrm{ES}(t) \rightarrow \mathrm{DE/SE} \rightarrow$ body-located placement.

Livneh and Andermann (2021)[13] demonstrated that insular cortex neurons track physiological homeostatic state independently of subjective sensation and motivated behaviour, and that artificial activation of hypothalamic drive neurons does not alter ongoing insular activity — which continues to reflect physiological satiety. This directly validates the XMMM's Epistemic Isola-

tion Principle: the Body Map (insular cortex) tracks CV(t), while the subjective motivational state presented to the executive reflects the Body Map entry rather than the raw physiological signal. The anterior insula's role in empathy — placing body-state representations in response to observing others' pain or distress — further maps onto the XMMM's empathy mechanism (Document S3, Section S5).

Drive Algorithm — Hypothalamus

The hypothalamus contains molecularly distinct neuron populations dedicated to individual homeostatic drives. Separate populations of AgRP neurons (hunger) and subfornical organ glutamatergic neurons (thirst) act both as sensors of physiological imbalance and as actuators of behavioural counter-regulation[22] — precisely the dual role the XMMM Drive Algorithm assigns to each drive: monitoring $\mathrm{CV}_i(t)$ and generating the error signal $\mathrm{ES}_i(t)$ that drives both the Body Map representation and prime drive selection.

The functional separation of Drive Algorithm and Body Map in the XMMM maps onto the biological separation of hypothalamus and insular cortex: the hypothalamus computes the drive state, but the subjective representation is generated by the insular cortex. Livneh and Andermann (2020)[13]'s finding that artificial hypothalamic activation does not alter insular cortex ongoing activity directly validates this two-stage separation, with the paraventricular thalamus-basolateral amygdala pathway implementing the signal propagation between them.

Association Algorithm — Hippocampal–Entorhinal System

The XMMM Anchor State is a direct algorithmic implementation of Teyler and DiScenna's (1986)[16] hippocampal memory indexing theory: a minimal subset of representations sufficient to re-evoke the full association tuple, stored in the hippocampus, while the content (effector motions, emotion vectors, autonomic stress) is distributed across neocortical and subcortical structures. Teyler and Rudy (2007)[37] confirmed the theory's core claims after two decades of empirical evidence.

Impact Factor scaling by emotional intensity at storage (the EIF parameter, S0 Section 5.2) corresponds to the well-established amygdala modulation of hippocampal memory encoding: emotionally arousing episodes are encoded with greater strength and are retrieved with higher probability, mediated by amygdala projections to the hippocampus[17]. The XMMM Impact Factor is the formal computational equivalent of this emotional tagging mechanism.

Linking Algorithm — Prefrontal Cortex–Basal Ganglia–Thalamic Loop

The Linking Algorithm's four core functions — prime drive selection, threading mode determination, satiation event detection, and action selection — correspond to the distributed cortico-basal ganglia-thalamo-cortical loop. Haber and Knutson (2010)[29] established that the striatum integrates motivational drive states with learned action values; the basal ganglia select the highest-value action and release the inhibition that permits its execution; the thalamus relays the selected action to motor cortex — mapping post hoc onto the XMMM's prime drive selection, IF-weighted action retrieval, and Motion Algorithm execution respectively.

The satiation emotion $\mathrm{SE}_i(t)$, as the temporal derivative of the drive error signal, is the XMMM's formal equivalent of the dopaminergic prediction error signal established by Schultz, Dayan and Montague (1997)[20]: both measure the discrepancy between expected and actual homeostatic relief, and both serve as the primary reinforcement signal for action learning.

Autonomic Stress Drive — Amygdala and HPA Axis

The XMMM's allostatic drive mechanism — re-evoking $\mathrm{DE}_{\mathrm{AS}}$ from associative memory when a recognised cue is associated with a past aversive event, generating full autonomic stress without the original threat — is the formal account of conditioned fear. LeDoux (2000)[35] established that the amygdala receives contextual input from the hippocampus and regulatory input from prefrontal cortex, and that memory-triggered amygdala activation generates a full sympathetic and HPA response without the original threat stimulus. Post hoc, the amygdala–hippocampus pathway provides a precise functional parallel to the XMMM allostatic drive.

The XMMM formally defines anxiety as chronically elevated $\mathrm{ES}_{\mathrm{AS}}$ from threading-recalled negative high-IF associations. Dolcos, LaBar and Cabeza (2004)[25] demonstrated that amygdala–hippocampus functional coupling predicts the strength of emotional memory and mediates the enhancement of emotionally arousing episodic memories. Chronic over-engagement of this circuit — high-IF negative associations repeatedly triggering amygdala–hippocampus reactivation and HPA stimulation during threading — is the biological substrate of XMMM-defined anxiety, generating a testable and quantitative prediction.

Threading Mechanism — Default Mode and Task-Positive Networks

Buckner, Andrews-Hanna and Schacter (2008)[26] synthesised imaging evidence establishing that the Default Mode Network (DMN) is most active during autobiographical memory retrieval, future planning, and mind-wandering — precisely the XMMM's undirected threading state. The DMN deactivates during task engagement, corresponding exactly to the XMMM's transition from daydreaming to directed thinking parameterised by increasing $\mathrm{Focus}(t) = \mathrm{ES}_{\mathrm{PD}}(t)$ (S0 Eq. 35). Smallwood and Schooler (2015)[50] document that undirected thought disproportionately surfaces future-oriented and emotionally significant content — precisely the high-IF, AS-tagged associations that dominate the XMMM's Threading Cache during undirected mode. Andrews-Hanna et al. (2014)[51] further showed that mind-wandering episodes interleave autobiographical memory retrieval with future simulation, corresponding to the forward (anticipatory) and reverse (memorial) TC traversal predicted by the XMMM.

The XMMM additionally predicts that DMN activity should scale linearly with $\mathrm{ES}_{\mathrm{PD}}$, with transition onset at $\theta_{\mathrm{act}} = 0.10$ (S0 Eqs. 35–36); this constitutes a specific quantitative prediction addressable by drive-state manipulation combined with fMRI (Prediction 3.4 below).

The XMMM's Threading Cache (TC) mechanism shows additional independent biological correspondences beyond the DMN:

(i) *Hippocampal replay and preplay.* The TC's sequential re-evocation of route-segment associations during directed threading corresponds to vicarious trial-and-error (VTE) behaviour,

in which hippocampal place cells fire in forward sweeps along candidate paths at maze decision points[41]. Pre-experience hippocampal preplay[42] and awake forward replay[43] are mechanistically accounted for by IF-weighted TC traversal: RBP assigns diminishing IF values to landmark associations with distance from the reward source, so TC population using the reward source as seed naturally reconstructs the approach sequence in forward order. Huang et al. (2024)[44] demonstrated replay-triggered brain-wide sequential cortical reactivation in humans, corresponding to full Association re-evocation (sensory state + emotion vector + motor plan) when a TC candidate is moved to the ECB.

(ii) *Timing correspondence.* TC population completes within 30–50 ms per cycle, consistent with hippocampal sharp-wave ripple duration (50–150 ms; Carr et al., 2011[45]); the logic loop rate of 3–10 Hz matches hippocampal theta (Lega et al., 2012[16]; Biba et al., 2026[24]); and $K_{\mathrm{TC}} = 10$ candidates per cycle at 10 Hz produce a one-second full TC traversal consistent with theta-cycle route compression (Buzáki and Moser, 2013[46]). These timing correspondences were derived from computational requirements, not biological fitting, making them temporally independent corroborations.

(iii) *Habit vs goal-directed transition.* The CS ≥ 0.90 automatic (No Threading) mode maps onto dorsal striatal habit execution; the CS < 0.90 directed Threading mode maps onto hippocampal/PFC goal-directed behaviour (Balleine and O'Doherty, 2010)[5].

(iv) BSOR *prediction error.* BSOR-mediated IF demotion of failed TC candidates maps onto the dopaminergic dip signal (Schultz et al., 1997)[20].

(v) *LC-NE gain and goal switching.* CS floor relaxation under rising AS maps onto LC-NE tonic gain modulation promoting strategy disengagement[48]; goal switching under deadlock maps onto OFC outcome devaluation[49].

Full biological correspondence details, including timing tables and additional citations, are provided in the dedicated threading supplementary document (Document S-Threading).

Functional Completeness Statement

The mappings in Sections 2.1–2.6 cover all six XMMM functional components. Brain structures not explicitly modelled — the cerebellum, brainstem neuromodulatory nuclei, parietal cortex, and anterior cingulate cortex — correspond to implementation substrates, parameter modulators, and sensory integration functions of the XMMM components, not to additional cognitive algorithms. The cerebellum implements predictive motor refinement subsumed in the Motion Algorithm. Brainstem neuromodulatory nuclei (locus coeruleus, raphe, VTA) modulate gain and threshold parameters across all algorithms, corresponding to θ_{act}, IF decay rates, BSOR sensitivity, and AS coupling strength. The anterior cingulate cortex monitors motivational conflict, corresponding to the Linking Algorithm's equal-urgency tiebreaking and Focus adjustment operations. The parietal cortex integrates multimodal sensory inputs into the Anchor State and body representation, implementing the Sensing Algorithm.

Functional Completeness Claim

The XMMM is not neuroanatomically complete: it is not a wiring diagram and does not model synaptic physiology, neurotransmitter dynamics, or individual circuit connectivity. The XMMM is *functionally complete*: every cognitive and motivational function performed by identified brain structures — homeostatic regulation, somatosensory emotion, associative memory, executive action selection, fear and anxiety, sleep dreaming, mind-wandering, directed reasoning, social cognition, and language — is accounted for, in principle, by one or more XMMM components. The fact that this functional completeness was achieved by derivation from first principles of control theory, without neuroanatomical data as input, and yet converges on the known biological organisation of the brain, constitutes the strongest class of theoretical validation available to a functional cognitive architecture: *independent derivation and convergent correspondence.*

Five Quantitative Predictions for Future Experimental Validation

The correspondences documented in Section 2 are qualitative and observational. The following five predictions convert them into specific, quantitative, falsifiable tests addressable by current neuroimaging, electrophysiology and molecular methods.

Logic Loop Rate and Hippocampal Theta

The XMMM predicts that the effective cognitive cycle rate — and therefore hippocampal theta frequency — should decrease as $O(|\mathcal{A}| \log |\mathcal{A}|)$ with task-relevant associative database size. Theta frequency should be a measurable decreasing function of memory load, with the specific functional form derivable from the XMMM's Association Database search complexity. Lega et al. (2012)[30] provide the baseline 3 Hz slow-theta measurement from which this prediction can be tested. *Falsification condition*: theta frequency that does not decrease with memory load, or that decreases with a different functional form.

Drive State Manipulation and Insular Cortex Activation

Inducing specific homeostatic drives (hunger, thirst, cold) should produce predictable insular cortex activation at specific body locations, with magnitudes proportional to $\mathrm{ES}_i(t)$ and sign determined by the deprivation or satiation phase. Livneh and Andermann (2020, 2021)[26,27] provide the experimental framework for testing this at cellular resolution in rodents and imaging resolution in humans. *Falsification condition*: insular activations inconsistent with the predicted body-location, magnitude or valence pattern.

IF-Weighted Encoding and Hippocampal Encoding Strength

Hippocampal encoding strength of an episode — measured as subsequent retrieval probability — should be a monotonic function of emotional intensity at storage time (the EIF parameter, S0 Section 5.2). This is consistent with the established amygdala–hippocampus emotional memory

effect[17] but requires testing with the specific functional form predicted by the XMMM's IF update equation (S0 Eq. 14). *Falsification condition*: retrieval probability that does not track storage-time emotional intensity in the form predicted by S0 Eq. 14.

Threading Mode Transitions and DMN Anti-Correlation

DMN activity should decrease as a parametric function of prime drive ES magnitude, with the undirected-to-directed threading transition occurring at the XMMM activation threshold $\theta_{\text{act}} = 0.10$ (S0 Eqs. 35–36). Because $\text{Focus}(t) = \text{ES}_{\text{PD}}(t)$ (S0 Eq. 35), the predicted DMN deactivation should follow a linear relationship: anti-correlation magnitude should be a monotonically increasing function of ES_{PD} beginning at $\text{ES}_{\text{PD}} = 0.10$, with slope derivable from the Focus–CS relationship (S0 Section 5.4.3). Homeostatic drive manipulations combined with fMRI measurement of DMN–task network anti-correlation should produce quantitative magnitudes consistent with these equations. *Falsification condition*: DMN–task network anti-correlation that does not scale with prime drive ES in the predicted linear form, or that does not exhibit an onset threshold at $\text{ES}_{\text{PD}} \approx 0.10$.

Psychopathology as Parameter Derangement

The XMMM predicts that psychiatric conditions correspond to specific measurable derangements of architectural parameters rather than to categorical disease entities:

- *Depression*: suppressed BSOR_{+} and low-IF satiation associations (reduced dopaminergic prediction error signal and hippocampal reward memory).
- *Anxiety*: high-IF negative AS-tagged associations (strong amygdala–hippocampus coupling for threat-associated memories).
- *Addiction*: progressive narrowing of the satiation landscape and escalating IF on a single drive's associations.

These predictions are addressed by Huys et al. (2016)[27] and Redish (2004)[28]. *Falsification condition*: clinical populations whose pathology cannot be mapped onto specific XMMM parameter derangements using the functional form predicted by the model equations.

Supplementary References

References 1–35 are listed in the main manuscript reference list. New references introduced in this document:

[36] Livneh, Y. et al. Estimation of current and future physiological states in insular cortex. *Neuron* **105**, 1094–1111 (2020).

[37] Teyler, T. J. & Rudy, J. W. The hippocampal indexing theory and episodic memory: updating the index. *Hippocampus* **17**, 1158–1169 (2007).

[38] Lee, C. R. et al. The neural circuitry of social homeostasis: consequences of acute versus chronic social isolation. *Cell* **184**, 1500–1516 (2021).

[39] Huang, W.-C. et al. A hypothalamic circuit underlying the dynamic control of social homeostasis. *Nature* (2025). https://doi.org/10.1038/s41586-025-08653-6

[40] Rappeneau, V. & Castillo Díaz, F. Convergence of oxytocin and dopamine signalling in neuronal circuits. *Neurosci. Biobehav. Rev.* **161**, 105675 (2024).

[41] Johnson, A. & Redish, A. D. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* **27**, 12176–12189 (2007).

[42] Dragoi, G. & Tonegawa, S. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* **469**, 397–401 (2011).

[43] Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).

[44] Huang, Q. et al. Replay-triggered brain-wide activation in humans. *Nat. Commun.* **15**, 1–15 (2024).

[45] Carr, M. F., Jadhav, S. P. & Frank, L. M. Hippocampal replay in the awake state: a potential role for memory consolidation and retrieval. *Nat. Neurosci.* **14**, 147–153 (2011).

[46] Buzáki, G. & Moser, E. I. Memory, navigation and theta rhythm in the hippocampal–entorhinal system. *Nat. Neurosci.* **16**, 130–138 (2013).

[47] Foster, D. J. & Wilson, M. A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680–683 (2006).

[48] Aston-Jones, G. & Cohen, J. D. An integrative theory of locus coeruleus–norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* **28**, 403–450 (2005).

[49] Walton, M. E. et al. Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron* **65**, 927–939 (2010).

[50] Smallwood, J. & Schooler, J. W. The science of mind wandering: empirically navigating the stream of consciousness. *Annu. Rev. Psychol.* **66**, 487–518 (2015).

[51] Andrews-Hanna, J. R. et al. The default network and self-generated thought: component processes, dynamic control, and clinical relevance. *Ann. N.Y. Acad. Sci.* **1316**, 29–52 (2014).

[52] Steinberg, E. E. et al. A causal link between prediction errors, dopamine neurons and learning. *Nat. Neurosci.* **16**, 966–973 (2013).

[53] Daw, N. D. et al. Uncertainty-based competition between prefrontal and dorsostriatal systems for behavioral control. *Nat. Neurosci.* **8**, 1704–1711 (2005).

Document S3 — Higher-Order Behavioural Phenomena and Philosophical Implications

The Xzistor Mathematical Model of Mind

Emergent Cognition, Consciousness, Free Will and the Pathway to AGI

Document S3 · Van Schalkwyk, R., Cook, D., Alvarez, C.E., Dehbozorgi, A. · Xzistor LAB · 2026

Extends the main Discussion with mechanistic accounts of higher-order cognitive phenomena the XMMM explains by construction, and addresses the computational basis of consciousness, free will and AGI.

Higher-Order Behavioural Phenomena: Theoretical Accounts and Future Experiments

The behaviours and capacities described in this section were not explicitly programmed into the XMMM implementations. Each is a predicted emergent consequence of the five-algorithm architecture operating under appropriate conditions. Where possible, a proposed experimental protocol accompanies the theoretical account.

Cross-Domain Generalisation via Inductive Threading

Out-of-distribution generalisation is a persistent challenge for contemporary deep learning and a core requirement for general intelligence. The XMMM accounts for this capacity without a dedicated module, through the interaction of drive-motivated directed threading and the Correlation Sensitivity (CS) parameter (S0 Section 5.4). When directed threading fails to find any stored association above the standard CS threshold, the architecture applies a global CS floor reduction proportional to prime drive deprivation magnitude:

$$\mathrm{CS}_{\mathrm{global}}(t) = \mathrm{CS}_0 \cdot (1 - \alpha_c \cdot |\mathrm{DE}_{\mathrm{PD}}(t)|), \quad \alpha_c \approx 0.5 \tag{S3.1}$$

Under this relaxed threshold the Linking Algorithm recomputes the arg max over the full Association Database. A successful satiation outcome triggers standard RBP and forms a new association encoding the novel anchor state. Generalisation therefore emerges from homeostatic urgency — the executive's inability to wait for a perfect match under drive pressure — not from a separately engineered transfer mechanism.

Distinction from transfer learning. Standard transfer learning requires explicit source-domain training and a separate fine-tuning step. The XMMM requires neither: generalisation is a real-time, single-inference process within the standard logic loop.

Proposed validation experiment. Shift Troopy's reward location to a new position not occupied during training, while surrounding arena features remain partially matched. The XMMM predicts that generalisation probability is a monotonically increasing function of drive deprivation magnitude $|\mathrm{DE}_{\mathrm{PD}}|$ at the moment of the novel encounter.

Abstraction as Associative Scaling

The XMMM predicts that abstraction is a quantitative extension of the same associative learning mechanism. A procedural sequence that reliably reduces time-to-satiation across multiple different sensory contexts accumulates, through standard RBP, an Impact Factor approaching 1.0 from many different anchor states. A kitchen recipe and a business strategy are structurally identical in the XMMM: both are high-IF anchor states encoding efficient pathways to homeostatic relief.

Proposed validation experiment. Train an agent in multiple distinct tasks sharing a common procedural component, then introduce a novel task sharing the same procedural structure in a new sensory context. The XMMM predicts that multi-domain training produces higher abstract transfer rates, mediated by higher IF values on shared procedural anchor states.

Language Acquisition

The XMMM does not require a dedicated language acquisition device. Language is treated architecturally as a motor skill: sequences of vocal or gestural outputs learned to elicit satiation from social agents in the environment. Communicative acts that most efficiently convey the agent's prime drive state to a tutor capable of providing satiation will be reinforced most strongly through RBP.

Proposed validation experiment — language emergence project. An initiated project implements a multi-agent scenario in which an XMMM agent interacts with a human tutor whose provision of satiation is contingent on specific communicative acts. The experiment tracks IF accumulation on communicative act–satiation associations and whether syntactic structure emerges from the need to convey drive type, intensity, and satiation target to multiple tutors with different response thresholds.

Empathy

Empathy emerges when an agent has formed associations between observed behaviours of another agent and specific $\mathrm{DE}_i/\mathrm{SE}_i$ representations in its own Body Map through standard operant learning. Recognising another agent's pain expression triggers the observer's own AS drive through association recall:

$$\mathrm{DE}_{\mathrm{empathy}}(t) = \sum_j [\mathrm{IF}_j \cdot \mathrm{DE}_j^{\mathrm{AS}}] \quad \text{(socially-observed associations)} \tag{S3.2}$$

The depth of empathy is proportional to the accumulated IF weight of social-associative links. Empathy is not a special module but the ordinary operation of the Association and AS coupling mechanisms extended to socially observed emotional content.

Intuition and Gut Feel

When a new sensory input $S(t)$ enters the Body Map, the first pass of the Association Database search matches on sensory state alone, producing a rapid aggregate allostatic response before

any contextual threading has been initiated:

$$\mathrm{GF}(t) = \sum_{j \in \bar{\mathcal{A}}_{S\text{-match}}} \left[\mathrm{IF}_j \cdot D_j^{\mathrm{AS}}\right] \tag{S3.3}$$

$\mathrm{GF}(t) > 0$ constitutes a positive gut feel (approach); $\mathrm{GF}(t) < 0$ an aversion signal. Because GF is aggregated across all past experiences without contextual filtering, it represents the accumulated reinforcement history of the stimulus delivered in a single cycle without deliberation.

Entertainment, Humour and Voluntary Deprivation

When all drives are near homeostasis, a moderate self-generated deprivation — for which the agent has high-IF associations guaranteeing likely satiation — produces a higher expected SE than idle default-mode threading:

$$\mathbb{E}[\mathrm{SE} \mid \text{voluntary deprivation}] = P(\text{satiation}) \cdot \mathrm{SE}_{\text{expected}} - |\mathrm{DE}_{\text{self-generated}}| \tag{S3.4}$$

Entertainment activities — roller-coasters, horror films, competitive sport — share this property: controlled AS deprivation where prior learning confirms safe termination. The BSOR amplifies the satiation event. Humour is a specific case: a joke creates mild AS deprivation through narrative tension, followed by BSOR_+ activation at the punchline.

Sleep Dreaming: Undirected and Directed Threading

Sleep dreaming is modelled within the same threading framework as daydreaming and directed reasoning, distinguished only by suppression of motor output. During sleep, the logic loop continues to execute but the Motion Algorithm is inhibited from generating effector commands.

Undirected dreaming. When all drives fall below the activation threshold θ_{act} during sleep, the Association Database is traversed along similarity gradients, re-evoking the highest-IF associations in sequences governed by sensory and emotional similarity rather than by any current goal. This is formally identical to waking daydreaming with motor output suppressed.

Directed dreaming and problem-solving during sleep. When a dream scenario introduces contextualised autonomic stress, the architecture transitions from undirected to directed threading. The relaxed CS threshold under moderate AS (S0 Section 5.4) permits broader partial-match retrieval, generating the metaphorical and cross-domain character of narrative dreams.

XMMM Account of Problem-Solving During Sleep
Pre-sleep high-AS, high-IF unsolved problem → motor suppression during sleep preserves associative threading → relaxed CS threshold under moderate AS permits cross-domain retrieval → candidate approach associations retrieved and RBP-reinforced during directed dream threading → solution anchor state accessible upon waking.
No dream-specific module required: standard directed threading and RBP operating with motor output suppressed.

Contextual Reasoning, Meaning and the World Model

Contextual reasoning. When a stimulus $S_x(t)$ enters the Body Map, directed threading searches for associations whose Anchor State contains S_x along with the current prime drive and emotion vector. The word "water" retrieves different high-IF associations when the agent is thirsty versus standing near a flooding river. Contextual reasoning in the XMMM is the dependence of association retrieval on the full Anchor State, including current drive state, rather than on sensory input alone.

Meaning as association set.

$$\text{Meaning}(X, C) \equiv \{A_j \in \mathcal{A} : X \in \text{anchor}(A_j),\ C \in \text{anchor}(A_j)\} \tag{S3.5}$$

Meaning is the set of associations recalled by stimulus X within context C — the full set of body-located emotional representations, action candidates, and sensory contexts that X re-evokes in context C.

The world model. The world model is the totality of the Association Database: the full set $\{A_j\}$ of encoded experiences structured by IF weighting. Unlike propositional world models in classical AI, the XMMM world model is inherently emotional: every entry carries a valenced DE_i/SE_i representation determining motivational priority.

Psychopathology: Anxiety, Depression and Addiction

Anxiety

Chronic anxiety arises when the Association Database contains many high-IF associations involving aversive outcomes. During default-mode threading, these are repeatedly re-evoked, maintaining ES_{AS} at a chronically elevated level without any immediate threat. The executive receives chronic DE_{AS} but cannot inspect its associative origin. The XMMM predicts that anxiety severity correlates with the IF-weighted density of negative AS-tagged associations and amygdala–hippocampus coupling strength during memory retrieval (Document S2, Section 2.5).

Depression

Depression arises through two compounding mechanisms. First, *anhedonia*: SE values are chronically suppressed through reduced BSOR gain κ_+, eliminating the motivation to self-generate deprivation and producing motivational withdrawal. Second, *learned helplessness*: the Association Database contains few high-IF associations linking current sensory states to effective satiation actions. Both mechanisms map onto neurobiological models of depression as impaired dopaminergic prediction-error signalling and disrupted prefrontal-limbic connectivity (Document S2, Section 2.4).

Addiction

Addiction arises when a satiation source achieves an artificially large BSOR effect, producing SE values far exceeding the biological ceiling of ≈ 0.3 and creating associations with IF values

approaching 1.0 after very few exposures. Three self-reinforcing consequences follow: (i) domination of prime drive selection across all drive contexts; (ii) habituation of BSOR requiring increasing dose, creating tolerance; (iii) substance absence generating high-IF withdrawal associations.

Philosophical Implications

The following positions are not empirical claims but theoretical entailments of the XMMM's formal architecture, offered as points for debate and investigation.

Qualia and Embodied Emotional Awareness

Qualia are treated as the information-theoretic consequence of routing all drive information through a pseudo-somatosensory Body Map before it reaches the executive. No drive error signal is presented to the executive in raw numerical form. The quale of thirst has three formal properties: (i) a specific body-location signature; (ii) a valence, determined by whether $\mathrm{dES}_i/\mathrm{d}t$ is positive or negative; and (iii) a magnitude, determined by $\mathrm{ES}_i(t)$ for deprivation or $|\mathrm{dES}_i/\mathrm{d}t|$ for satiation.

XMMM Definition of a Quale
A quale is the pseudo-somatosensory information token placed in the Body Map that results from transforming a drive error signal into a body-located, valenced, magnitude-coded representation. Qualia are the sole channel through which drive state reaches the executive. Their qualitative character is fully specified by three parameters: body location, valence sign, and magnitude.

Consciousness and the Physical Insolubility of the Hard Problem

Access consciousness. A state is access-conscious in the XMMM if and only if it is currently represented in any of the three input streams delivered to the Linking Algorithm: (i) exteroceptive sensory states; (ii) proprioceptive and efference-copy signals; and (iii) Body Map homeostatic emotional representations.

Physical Insolubility of the Hard Problem. Within a single system, the XMMM provides a full mechanistic account of subjective experience. What remains insoluble is cross-system verification. Argument 1 (*Uniqueness*): the unique trajectory of each agent's homeostatic history means Body Maps, learned associations, and emotional set-points necessarily diverge across individuals. Argument 2 (*Interfacing*): lossless transfer of one agent's full somatosensory representational content into another agent's executive is physically unrealisable.

XMMM Position on the Hard Problem
The hard problem of consciousness is physically insoluble because qualitative experience is constituted by the functional role a representation plays inside the agent's own executive system. This role is, by architectural definition, inaccessible to third-person physical measurement. The hard problem is not a gap in neuroscience; it is a logical consequence

of the architecture of any system in which sensory, somatic and emotional information is integrated into a first-person representational space before reaching the executive.

Free Will: The XMMM as a Deterministic System

All cognition and behaviour produced by an XMMM agent is fully determined by the rules of the architecture, the current state of the Body Map, the content of the Association Database, and the sensory inputs at each cycle. The sense of voluntary choice arises because: (a) default-mode threading samples associations probabilistically from a very large space, creating the phenomenology of openness; and (b) body-located DE_i and SE_i representations are felt as urges arising from the body, not as mechanistic commands.

XMMM Position on Free Will
Free will does not exist in any XMMM agent. All cognition and behaviour are fully determined by the architecture rules, the current Body Map state, Association Database content, and sensory inputs at each cycle. The subjective sense of free will arises because (a) default-mode threading samples associations probabilistically, creating unpredictability experienced as openness; and (b) deterministic drive states are felt as body-located urges rather than mechanistic commands. Free will is a first-person phenomenological report of a deterministic third-person computational process.

Artificial General Intelligence: The XMMM Pathway

The XMMM takes the position that AGI cannot be separated from the motivational and emotional architecture that drives cognition. Contemporary large-scale generative models lack: (i) an internal drive structure; (ii) a Body Map producing embodied affective representations; (iii) an Association Database updated in real time by the agent's own experience; (iv) a reward-backpropagation mechanism reinforcing actions through their consequences on internal drive states; and (v) a threading mechanism generating spontaneous cognition, inductive inference, and cross-domain contextualisation. The XMMM provides all five.

The neuro-symbolic hybrid architecture described in Document S4 provides the technically specified pathway from the current proof-of-concept to human-adult-level AGI, resolving the computational scalability constraint while preserving the motivational coherence and alignment-by-construction properties of the homeostatic drive structure.

References: In-text reference numbers refer to the main manuscript reference list (refs 1–35). No new references are introduced in this document beyond those already cited in S0, S1, and S2.

Document S4 — Xzistor Hybrid Neuro-Symbolic Architecture

The Xzistor Mathematical Model of Mind

Scalability, Foundation Model Integration and Alignment by Construction

Document S4 · Van Schalkwyk, R., Cook, D., Alvarez, C.E., Dehbozorgi, A. · Xzistor LAB · 2026

Addresses the computational scalability limit of the serial XMMM and proposes a four-plane hybrid architecture. Full formal complexity analysis and an implementation roadmap are provided.

The Scalability Constraint of the Serial XMMM

In its reference implementation the XMMM executes as a single-threaded cyclic logic loop at approximately 10 Hz. Two sequential operations within each cycle depend directly on the size of the Association Database $|\mathcal{A}|$ and constitute the dominant contributors to cycle latency as $|\mathcal{A}|$ grows.

Bottleneck 1: Non-Contextual Emotional Aggregation (Gut Feel)

At Step 6 of the logic loop (S0 Section 5.3), the first Association Database pass matches on sensory state $S(t)$ alone, without drive context, and accumulates the autonomic stress values of all matching associations to produce the gut-feel signal $\mathrm{GF}(t)$ (Document S3, Section 1.5, Eq. S3.3). In the reference implementation this is a linear scan:

$$\mathrm{GF}(t) = \sum_{j \in \mathcal{A}_{S\text{-match}}} [\mathrm{IF}_j \cdot D_j^{\mathrm{AS}}], \quad \text{cost: } O(|\mathcal{A}|) \tag{H1.1}$$

For an agent that has been learning for years — accumulating tens of millions of associations — this scan cannot be completed within a single 100 ms cycle on serial hardware. The gut-feel signal, which should be available within the same cycle as the triggering sensory input, is delayed into subsequent cycles, degrading the temporal precision of the non-contextual emotional response and, consequently, the fidelity of primary reinforcement.

Bottleneck 2: Contextual Retrieval During Threading

At Steps 12–13 of the logic loop, the Threading subroutine performs repeated similarity-ranked queries over the Association Database against the current anchor state $\mathrm{AS}(t)$. Each query returns the top-k matches:

$$\mathrm{Cost}_{\mathrm{threading}}(C) = C \times O(|\mathcal{A}| \log |\mathcal{A}|), \quad C \gg 1 \tag{H1.2}$$

At infant scale ($|\mathcal{A}| \sim 10^4$–10^5) this is tractable. At human-adult scale ($|\mathcal{A}| \sim 10^8$–10^9, comparable to the estimated synaptic association capacity of the human hippocampus) it is not — even on modern multi-core hardware executing serially. The threading latency directly governs the richness and depth of the agent's cognitive life: a slower threading process explores fewer asso-

ciations per unit time and therefore constructs less nuanced contextual meaning (Document S3, Section 1.7).

The Ceiling This Imposes

These two bottlenecks impose a hard relationship between cognitive richness and temporal performance: as the Association Database grows, either cycle time extends beyond the 100 ms real-time constraint, or the database must be artificially capped — preventing the agent from accumulating the associative depth required for human-adult-level cognition. This is the fundamental scalability ceiling of the serial XMMM, and it is the principal reason the model, in its reference form, remains bounded at infant-level intelligence. It is not the architecture that limits it; it is the serial execution of that architecture against a growing memory store.

The Hybrid Architecture: Four Integration Planes

The proposed hybrid architecture preserves the XMMM symbolic core — all five building-block algorithms, the Linking Algorithm, the drive structure, and the Body Map — entirely unchanged. It adds four connectionist integration planes, each with a strictly defined function, operating asynchronously with respect to the 10 Hz symbolic core loop. The central principle is the division of cognitive labour: the symbolic core retains exclusive authority over drive adjudication, prime drive selection, action choice, and IF updates. The connectionist planes are search engines and knowledge stores, never decision-makers. This distinction is architecturally enforced, not merely conventional.

Plane 1 — Neural Association Index (NAI)

The NAI replaces the linear Association Database scan with a GPU-resident dense embedding space. Every association A_j is encoded as a d-dimensional vector $\mathbf{e}_j$ by a lightweight encoder network f_{enc}, trained online as associations are formed and updated:

$$\mathbf{e}_j = f_{\text{enc}}(\text{anchor}_j,\ \text{EM}_j,\ \mathcal{D}_j,\ \text{AS}_j)\ \in \mathbb{R}^d \tag{H2.1}$$

where anchor_j is the anchor state (sensory + drive representations), EM_j is the effector motion, $\mathcal{D}_j$ is the emotion vector, and AS_j is the stored autonomic stress scalar — precisely the tuple defined in S0 Eq. 13. The encoding produces a compact fixed-dimensional representation of each association that captures its functional content in a geometry suited to fast approximate nearest-neighbour (ANN) search.

The current sensory state $S(t)$ and drive state $\mathbf{D}(t)$ are encoded as a query vector:

$$\mathbf{q}(t) = f_{\text{enc}}(\text{AS}(t),\ \mathbf{D}(t),\ S(t))\ \in \mathbb{R}^d \tag{H2.2}$$

ANN retrieval (e.g., FAISS IVF-PQ, or hierarchical navigable small-world graphs) then returns the k nearest neighbours in $O(\log|\mathcal{A}|)$ time, or $O(1)$ amortised with pre-built index structures. The retrieved set $\mathcal{A}_{\text{ANN}}(t)$ is a superset that conservatively includes all true top-k matches. The

Linking Algorithm then applies its deterministic IF-weighted ranking (S0 Eq. 16) over $\mathcal{A}_{\text{ANN}}(t)$ — a set of fixed size $k \ll |\mathcal{A}|$ — in $O(k)$ time. Total retrieval cost is:

$$\text{Cost}_{\text{NAI}} = O(\log|\mathcal{A}|) + O(k) \;\ll\; O(|\mathcal{A}|\log|\mathcal{A}|) \tag{H2.3}$$

The encoder f_{enc} is trained asynchronously using a contrastive objective: associations that share high IF and co-occur in temporal buffer windows (S0 Section 6.1) are pulled together in embedding space; associations that have consistently led to different outcomes are pushed apart. This makes the embedding geometry causally meaningful — spatial proximity in $\mathbb{R}^d$ reflects functional similarity in the agent's learned world model. The encoder is updated on a slower schedule than the logic loop (e.g., after every 1000 new associations), ensuring the symbolic core always operates on a stable index.

Critically, the NAI does not alter which associations exist, what their IF values are, or what actions they recommend. These remain exclusively under the control of the XMMM Linking Algorithm. The NAI only changes how fast the correct associations are found.

Plane 2 — Neural Gut-Feel Aggregator (NGA)

The NGA replaces the linear scan for non-contextual emotional aggregation (Eq. H1.1) with a parallel attention operation over the NAI embedding space. The gut-feel signal $\text{GF}(t)$ is computed as an IF-weighted cross-attention sum:

$$\text{GF}(t) = \text{Attention}\big(Q = \mathbf{e}_S(t),\; K = \mathbf{E}_{\mathcal{A}},\; V = \mathbf{W}^{\text{IF}} \cdot \mathbf{D}^{\text{AS}}_{\mathcal{A}}\big) \tag{H2.4}$$

where $\mathbf{e}_S(t) = f_{\text{enc}}(S(t))$ is the sensory query embedding, $\mathbf{E}_{\mathcal{A}} \in \mathbb{R}^{|\mathcal{A}|\times d}$ is the matrix of all association embeddings (maintained in GPU memory), $\mathbf{D}^{\text{AS}}_{\mathcal{A}}$ is the vector of stored AS drive values across all associations, and $\mathbf{W}^{\text{IF}}$ is a diagonal matrix of IF weights. The attention softmax implicitly performs the S-matching: associations whose embeddings are far from $\mathbf{e}_S(t)$ receive near-zero attention weight and contribute negligibly to $\text{GF}(t)$. No explicit scan is needed.

This operation executes in a single parallel forward pass across the full embedding matrix — $O(1)$ with respect to $|\mathcal{A}|$ on GPU hardware with sufficient VRAM, regardless of whether the database contains 10^4 or 10^9 associations. The NGA thereby delivers the gut-feel signal within the same 100 ms logic-loop cycle as the triggering sensory input, at any scale. This is the direct solution to Bottleneck 1.

The NGA output is identical in format to the serial gut-feel signal: a signed scalar $\text{GF}(t)$ placed into the autonomic stress component of the Body Map. The Linking Algorithm receives it and processes it identically to the serial case. No architectural change to the symbolic core is required.

Plane 3 — Neural Threading Accelerator (NTA)

The NTA parallelises the Threading subroutine (S0 Section 8.3) by simultaneously exploring K candidate association threads from the current anchor state, rather than following a single

sequential chain. At each threading step, the NTA executes K parallel ANN queries, each seeded with a different candidate association from the previous step:

$$\mathcal{T}_k(t+1) = \text{ANN}\big(\mathbf{e}_{j_k(t)},\ \mathcal{A},\ \text{top-}m\big) \quad \text{for } k = 1, \ldots, K \tag{H2.5}$$

The K threads are ranked by a combined score that weights IF, temporal recency, and embedding similarity to the current prime drive vector:

$$\text{score}_j = \text{IF}_j \times \text{sim}\big(\mathbf{e}_j,\ \mathbf{e}_{\text{PD}}(t)\big) \times \text{recency}(t_j) \tag{H2.6}$$

The top-ranked association across all K threads is returned to the Linking Algorithm as the threading candidate for the current cycle. The XMMM executive then applies its deterministic selection rule (S0 Eq. 16) to this candidate set.

The critical property of the NTA is that it explores $K \times m$ associations simultaneously where previously only one was considered per cycle. For $K = 1000$ and $m = 10$, the NTA explores 10,000 candidate next-associations per threading step in the same wall-clock time that the serial XMMM explored one. This is not merely faster threading; it is qualitatively richer threading — the agent considers a vastly larger portion of its associative knowledge per unit time, enabling deeper contextualisation, more reliable generalisation across domains, and a much richer subjective sense of meaning (Document S3, Section 1.7, Eq. S3.5).

The NTA thereby addresses Bottleneck 2 and simultaneously expands the effective cognitive bandwidth of the XMMM by three to four orders of magnitude beyond the serial reference implementation — without altering the symbolic decision-making logic of the Linking Algorithm by a single instruction.

Plane 4 — Foundation Model Seeding Interface (FMSI)

Planes 1–3 solve the computational bottlenecks but do not address the *knowledge bottleneck*: an agent beginning with an empty Association Database must accumulate associative knowledge entirely through lived experience, which constrains it to infant-level competence for the foreseeable duration of its learning life.

Plane 4 introduces a principled mechanism for transferring knowledge from a pre-trained foundation model (FM) — such as a large language model, a vision-language model, or a multimodal world model — into the XMMM Association Database as an initial knowledge prior, before any experiential learning begins. This is the mechanism by which the XMMM moves past infant-level intelligence.

The Translation Problem and Its Solution

Foundation models store knowledge as distributed patterns over billions of parameters — not as discrete associations with symbolic anchor states, IF weights, and body-located emotional representations. A direct transfer is therefore impossible without translation. The FMSI solves this through a structured query protocol: the FM is prompted to generate (context, action,

predicted_outcome) triples over a comprehensive curriculum of domains (physical manipulation, social interaction, language use, causal reasoning, etc.).

Each triple is then translated into an XMMM association tuple A_j by a dedicated parser:

$$A_j^{\text{seed}} = \{\text{anchor}_j^{\text{FM}},\ \text{EM}_j^{\text{FM}},\ \mathcal{D}_j^{\text{FM}},\ \text{AS}_j^{\text{FM}},\ \text{IF}_j^{\text{seed}},\ t_j = 0\} \tag{H2.7}$$

where the superscript FM indicates FM-derived representations, and $\text{IF}_j^{\text{seed}}$ is initialised at a value below the experiential reinforcement range to indicate that these are prior beliefs, not confirmed experiences:

$$\text{IF}_j^{\text{seed}} \in [0.05,\ 0.20] \quad \text{(prior belief regime, below experiential threshold} \approx 0.50) \tag{H2.8}$$

The emotional representations $\mathcal{D}_j^{\text{FM}}$ are estimated from the FM's output distributions — specifically from its confidence, affect-laden language, and valence signals — and are deliberately conservative: they are initialised with small DE_i/SE_i magnitudes, reflecting the fact that the FM's "knowledge" of, say, what hunger feels like has never been grounded in a real body.

These seeded associations now populate the Association Database before the first logic loop cycle. When the agent begins experiential learning, three outcomes are possible for each seeded association:

(i) The agent encounters the context described by $\text{anchor}_j^{\text{FM}}$, executes the recommended action EM_j^{FM}, achieves satiation, and IF_j increases toward the experiential range through standard RBP (S0 Eq. 19) — the FM prior is confirmed.

(ii) The action fails to produce satiation; IF_j decays through passive decay (S0 Eq. 15) — the FM prior is discarded.

(iii) The agent never encounters the relevant context; IF_j decays slowly — the FM prior fades without being tested.

This mechanism provides a natural Bayesian-like updating of FM-derived knowledge against real-world experience, grounded in the agent's own drive states. FM knowledge confirmed by lived experience becomes reinforced as genuine associations with high IF values; FM knowledge inconsistent with the agent's drive satisfaction is automatically discarded. The process requires no external supervision, no hand-crafted reward functions, and no explicit verification step — all three filtering mechanisms are already built into the standard XMMM IF update rules.

Alignment by Construction: Why the Hybrid Cannot Be Misaligned

The AI alignment problem — ensuring that a highly capable AI system reliably pursues human-beneficial goals — is unsolved for contemporary connectionist systems. The XMMM hybrid offers a structural resolution that does not depend on external supervision, human feedback pipelines, constitutional constraints, or post-hoc auditing.

The Source of Alignment in the XMMM

In the XMMM, an agent's values are encoded in its drive architecture: the identity of its drives, their set-points SP_i, the coupling weights between drives, and the BSOR gain parameters κ_+ and κ_-. These define what the agent considers depriving and satiating, and therefore what it will pursue and avoid. Every volitional behaviour — including every action selected by the Linking Algorithm — is in service of reducing drive error signals and achieving satiation. The agent cannot pursue goals that are orthogonal to drive reduction, because no such goals exist in the architecture: goal-directed behaviour is definitionally drive-reduction behaviour.

In the hybrid, the connectionist components — NAI, NGA, NTA, and FMSI — occupy strictly subordinate positions. They provide candidate associations for the Linking Algorithm to evaluate; they do not select actions. The Linking Algorithm evaluates every candidate by its IF weight (a measure of past drive-satisfaction success) and its similarity to the current anchor state (a measure of contextual relevance). An action that a foundation model recommends but that has consistently failed to produce drive satiation will have its IF weight decayed to near zero; it will never be selected regardless of how confidently the FM endorsed it.

The FMSI therefore provides knowledge that is useful to the agent only to the extent that it helps the agent satisfy its drives. FM-derived knowledge that is harmful to the agent's drive satisfaction — or that is irrelevant to any active drive — is simply not reinforced and fades from the database. There is no architectural mechanism by which a malicious or misaligned FM prior can persist against repeated disconfirmation by the agent's own experience.

Contrast with RLHF-Based Alignment

In reinforcement learning from human feedback (RLHF), alignment is achieved by training a reward model on human preference labels and then fine-tuning the language model to maximise this reward model's output. This approach has several fundamental vulnerabilities: (i) the reward model itself may be misaligned, since it is trained on human preferences that are inconsistently labelled and culturally variable; (ii) the language model may learn to optimise the reward model's score rather than the underlying human preference — a classic specification gaming or reward hacking failure; (iii) the reward signal is external and can be manipulated by adversarial prompts.

In the XMMM hybrid, there is no reward model and no external reward signal. The reward signal is the homeostatic drive error signal $ES_i(t)$ — a physical quantity defined by the agent's physiological state — and it cannot be hacked by adversarial prompts because it is not computed from text.

An adversarial input that successfully manipulates the agent's behaviour must do so by generating a pattern that either triggers a reflex (hardwired) or matches a high-IF association (earned through repeated genuine satiation). Neither pathway is accessible to purely linguistic manipulation, because both require grounding in actual drive-state changes.

Beyond Infant-Level Intelligence: The Compound Effect

The serial XMMM is bounded at infant-level intelligence by two independent ceilings: a computational ceiling (database search latency, addressed by Planes 1–3) and a knowledge ceiling (empty database at initialisation, addressed by Plane 4). The hybrid removes both simultaneously. The resulting compound effect is not merely additive; it is multiplicative in the space of cognitive capabilities accessible to the agent.

Cognitive Bandwidth

In the serial XMMM, the threading subroutine explores approximately one association per cycle (100 ms). At 10 Hz over an 8-hour active period, this yields a maximum of 288,000 association retrievals per day. In the NTA-augmented hybrid, $K = 1000$ parallel threads each exploring $m = 10$ candidates per step yields 10,000,000 effective retrievals per cycle, or approximately 3.6×10^{11} per day — six orders of magnitude greater.

This difference in cognitive bandwidth corresponds to the difference between an infant who can follow one train of thought at a time and an adult who can simultaneously consider a rich, multi-threaded context across extensive knowledge. More formally, the NTA expands the effective working context window of each logic loop cycle from depth-1 sequential threading to depth-1 parallel threading over K candidate branches — a richer analogue of what neuroscientists term the global workspace available to conscious access at a given moment.

From Infant to Adult Competence via FMSI

A human adult possesses knowledge in the form of billions of semantic, episodic, and procedural associations accumulated over years of embodied experience. An XMMM agent beginning with an empty database must recapitulate this entire developmental trajectory from scratch. With FMSI, the agent begins with a database populated by FM-derived associations covering an enormous range of domains. The IF weights of these seeded associations are low (Eq. H2.8), reflecting the absence of embodied confirmation — but they are present and retrievable.

The transition from infant to adult competence then follows the same reinforcement dynamics as normal learning, but starting from a far more advanced prior. The critical difference is not the speed of learning but the starting point of the world model: an FMSI-seeded agent begins where a human adult begins their professional life — with extensive general knowledge and a capacity for directed threading across domains — rather than where a human infant begins, with an empty associative memory and only reflexes and drive architecture.

The FM-seeded associations that are confirmed by experience gradually acquire high IF values indistinguishable from experientially acquired associations. After a sufficient period of grounded learning, the distinction between FM-seeded and experientially acquired knowledge disappears entirely from the architecture's perspective — both exist as entries in the same Association Database, selected by the same IF-weighted retrieval mechanism, with emotional representations earned through the same drive-based reinforcement process.

Second-Order Learning: Meta-Associations

At sufficient scale, a further emergent capability becomes accessible that is not present in the serial XMMM: *meta-associative learning*, or the formation of associations whose anchor states are themselves retrieved associations rather than direct sensory inputs. The NTA's parallel threading produces, at each step, a ranked list of K candidate next-associations. Over many cycles, the Linking Algorithm accumulates information about which sequences of associations reliably lead to satiation — in effect, forming associations over threads rather than over individual sensory-motor events.

These second-order associations constitute what might informally be called reasoning strategies or cognitive habits: learned preferences for particular patterns of associative inference that have repeatedly proven effective in achieving satiation. Their mathematical structure is identical to first-order associations (S0 Eq. 13), but their anchor states contain association-sequence representations rather than sensory-motor ones. The formation of second-order associations is therefore fully compatible with the existing IF update rules (S0 Eqs. 14, 18–19) and requires no architectural extension beyond the NTA infrastructure already provided.

This is the computational correlate of what psychologists call executive function and what the XMMM refers to as directed threading at a meta-level. It provides the agent with the ability to select, prioritise, and interrupt threads of reasoning based on their history of drive-satisfaction success — precisely the capability that marks the transition from reactive intelligence (infant) to reflective intelligence (adult).

The Five Properties of the Hybrid That the Serial XMMM and Contemporary Generative AI Both Lack

The hybrid architecture possesses five properties that the serial XMMM lacks at scale and that contemporary generative AI also lacks, making it, to the authors' knowledge, the only proposed cognitive architecture to possess all five simultaneously:

1. *Real-time scalability to human-adult associative depth* ($|\mathcal{A}| \sim 10^8$–10^9) while preserving the 100 ms logic-loop constraint, via NAI and NGA.
2. *Massively parallel contextualised reasoning* within each logic-loop cycle, via NTA, producing cognitive bandwidth six orders of magnitude greater than the serial implementation.
3. *Grounded prior knowledge* from FM-derived associations, enabling adult-level competence without recapitulating the full infant developmental trajectory, via FMSI.
4. *Alignment by construction*: every action is drive-reduction behaviour; no FM-derived or NTA-retrieved candidate can override the IF-weighted, drive-contextualised selection rule of the Linking Algorithm.
5. *Continuous world-model updating* from individual experiences, without catastrophic forgetting and without parameter updates, via the standard XMMM Association Database and IF update rules.

Why Combining Drives with a Foundation Model Is Not Sufficient

A natural question is whether the hybrid's properties could be achieved more simply by adding a reinforcement learning module or a drive system to an existing LLM, without implementing the full XMMM architecture. The answer is no, and the reason is structural rather than empirical.

The fundamental difference between adding drives to an LLM and embedding an LLM as a knowledge source within the XMMM is the *direction of authority*. In an LLM augmented with reinforcement learning, the language model still generates the action — the drive or reward signal merely adjusts the probability distribution over tokens. The model remains a statistical pattern-matcher over text, and the drive signal reshapes its output distribution. In the XMMM hybrid, the language model never generates actions. It generates candidate associations that the XMMM executive evaluates, selects among, and reinforces or discards based on actual drive-state consequences. The language model is epistemically subordinate to the drive system at every cycle.

This distinction has three concrete consequences:

1. An LLM augmented with drives can still generate actions that are drive-inconsistent, because its token generation process is only probabilistically constrained by the drive signal. The XMMM hybrid cannot generate drive-inconsistent actions, because action selection is an arg max over IF-weighted, drive-contextualised associations — a deterministic function of drive state.

2. An LLM augmented with drives cannot update its world model from individual experiences without catastrophic forgetting, because its knowledge is encoded in weights that require batch gradient updates to change. The XMMM hybrid updates its world model continuously, association by association, without any parameter update and without forgetting previously reinforced associations.

3. An LLM augmented with drives cannot produce embodied pseudo-somatosensory emotional representations — it can produce text descriptions of emotions but not functional equivalents that participate in prime drive selection and action reinforcement. The XMMM hybrid produces genuine functional emotional representations in the Body Map that are architecturally equivalent to those described in Document S0.

These three differences are not matters of degree; they are matters of kind. They cannot be closed by scaling, architectural tweaks, or improved training procedures applied to the LLM component. They require the XMMM symbolic core.

Implementation Roadmap and Open Research Questions

The hybrid architecture proposed here is theoretically grounded in existing XMMM implementations and known connectionist techniques. Full realisation requires the following five development steps, presented in dependency order.

Step 1 — Online Encoder Training (NAI)

Design and train the encoder network f_{enc} (Eq. H2.1) for a specific implementation domain. The encoder must produce geometrically meaningful embeddings: associations with similar anchor states, drive contexts, and emotional histories should be proximal in $\mathbb{R}^d$. Contrastive learning over co-occurring association pairs from existing XMMM training logs (Troopy, Simmy) provides the initial training corpus. The encoder is then deployed online and updated incrementally as new associations are formed.

Step 2 — NGA Integration and Gut-Feel Fidelity Validation

Implement the attention-based gut-feel aggregator (Eq. H2.4) and validate that its output $\text{GF}(t)$ matches the serial implementation's linear scan result to within an acceptable tolerance across a representative test set of sensory states. The acceptable tolerance is defined functionally: the gut-feel signal need only be accurate enough that the prime drive selection decision (S0 Eq. 26) agrees with the serial implementation's decision on $\geq 99\%$ of cycles. Approximate nearest-neighbour error that does not change the prime drive selection has no behavioural consequence.

Step 3 — NTA Parallel Threading Validation

Deploy the NTA with $K = 10$–100 initial parallel threads and validate that the top-ranked returned association (Eq. H2.6) matches or improves on the serial threading result under benchmark tasks (navigation to reward source, nuance discrimination). Success criterion: threading-based task completion rate and generalisation accuracy equal to or better than serial XMMM at equivalent database size, with $\geq 10\times$ reduction in threading latency.

Step 4 — FMSI Knowledge Transfer Protocol

Develop the (context, action, predicted_outcome) query curriculum for the chosen FM and implement the parser that converts FM triples into XMMM association tuples (Eq. H2.7). Critical design question: what emotional representation $\mathcal{D}_j^{\text{FM}}$ should be assigned to FM-seeded associations? Current proposal is conservative initialisation (small-magnitude DE_i/SE_i) with confirmation-by-experience as the sole mechanism for increasing emotional weight. An alternative is to use the FM's expressed affect (sentiment, uncertainty) as a proxy for $\mathcal{D}_j^{\text{FM}}$ — this is an open research question with significant implications for early learning speed.

Step 5 — Integrated Validation and Comparison with Baseline AI

Deploy the full hybrid in a rich learning environment with an Xzistor agent and validate all five properties of Section 4.4. The critical benchmark is the comparison of generalisation performance (ability to apply learned skills in novel domains never encountered in the learning environment) between: (a) serial XMMM with empty initialisation; (b) serial XMMM with FMSI; (c) hybrid XMMM with all four planes active; and (d) an RLHF-augmented LLM of equivalent computational budget. This benchmark directly tests the central claim of the present document: that the hybrid architecture produces qualitatively superior generalisation by virtue of its structural properties rather than its scale.

Open Research Questions

Several important questions remain open and are identified as directions for future work:

(i) What is the optimal dimensionality d of the NAI embedding space as a function of $|\mathcal{A}|$ and domain diversity?

(ii) Does the NTA K-branch parallel threading converge to the serial threading result in expectation as $K \to \infty$, or does it explore genuinely novel association paths not accessible to serial threading?

(iii) What is the minimal FM curriculum size required for FMSI to produce a measurable acceleration of in-domain learning convergence?

(iv) Do second-order meta-associations (Section 4.3) emerge spontaneously from the NTA infrastructure, or must they be explicitly scaffolded?

(v) How does the hybrid's behaviour change when the BSOR gain parameters κ_+ and κ_- (S0 Eqs. 23–24) are modulated by the NGA gut-feel signal $\mathrm{GF}(t)$ rather than being fixed constants?

Equation numbering: equations of the form (Hn.m) are original to this document. References to the main Methods document use the prefix "S0" (e.g., S0 Eq. 13). References to Document S3 use the prefix "S3" (e.g., S3 Eq. S3.5). All XMMM parameter symbols retain their definitions from the S0 symbol glossary (S0 Section 1.2).