

# Does XMMM Provide an Algorithm for Consciousness?

*A control-theoretic perspective for a multidisciplinary audience*

Rocco Van Schalkwyk | Xzistor LAB, Bristol, UK | May 2026

---

## 1. The Challenge

The assertion that there is no known algorithm for consciousness is commonly held across philosophy, neuroscience and AI. I disagree — and I want to explain why carefully, because the disagreement turns on a precise distinction that is easy to miss.

The challenge posed to me on X was direct: *"Doing sophisticated calculations doesn't imply being conscious."* I agree completely. The claim I am making is not that more computation produces consciousness. That would indeed be naive. My claim is different: that consciousness arises from a **specific functional architecture** — a minimal set of interoperating algorithms that, when run in an embodied agent, produce the conditions necessary for subjective experience. The Xzistor Mathematical Model of Mind (XMMM) specifies and implements that architecture. The distinction matters enormously, and most of the disagreement in this field collapses once it is clearly drawn.

## 2. What Consciousness Actually Requires

Before asking whether XMMM provides an algorithm for consciousness, we need to ask what consciousness actually requires at a functional level. I argue it requires four things — none of which is *computational complexity*:

- **A body — or a faithful simulation of one.** Not a brain in isolation. A system with internal states that deviate from set-points and generate valenced signals: good and bad, approach and avoid. Without a body, there is nothing it is *like* to be the system, because there is no state the system is trying to get away from or toward.
- **Homeostatic drives.** The body's deviations from set-points must be represented as emotionally valenced states — what I call Deprivation Emotions (DE) and Satiation Emotions (SE), placed in a Body Map at body-specific locations. These are not abstract reward signals. They are felt states. The hunger is *somewhere*, the pain is *somewhere*. This is the substrate of experience.
- **Sensorimotor integration in a closed loop.** The agent must sense, feel, act, and sense again — continuously. Consciousness is not a snapshot; it is the ongoing process of a system maintaining itself in a world that resists it.
- **A unified executive that acts on felt experience.** The XMMM's Linking Algorithm receives only body-located emotional representations — never raw physiological values — and selects the Prime Drive as the most urgent felt state. This is what it means to act on experience rather than data.

A system with all four properties has, I argue, the minimal sufficient conditions for consciousness. Not rich adult human consciousness — but consciousness nonetheless. The Wright Brothers' first flight was not a 747. It was flight.

### 3. The XMMM Algorithms for Consciousness

The XMMM comprises five substrate-independent functional algorithms — Sensing, Drive, Reflex, Association, and Motion — coordinated by a Linking Algorithm. Together they implement the four conditions above. Here is how each contributes:

XMMM Algorithm	Contribution to Consciousness
<b>Drive Algorithm</b>	Converts homeostatic deviations into body-located Deprivation and Satiation Emotions — the felt substrate of experience
<b>Sensing Algorithm</b>	Updates the Body Map continuously, creating the ongoing sensory dimension of subjective experience
<b>Reflex Algorithm (BSOR)</b>	Generates anticipatory autonomic stress — the prediction-error signal that makes unexpected outcomes emotionally salient
<b>Association Algorithm</b>	Encodes experience as emotionally-contextualised memory; enables the agent to act on what it has felt before
<b>Motion Algorithm</b>	Closes the loop — the agent acts on its felt experience and the world responds, sustaining the cycle
<b>Linking Algorithm</b>	Receives only body-located emotional representations; selects the Prime Drive; constitutes the unified executive 'point of view'

The critical property of this architecture is not complexity but **integration**. It is the continuous, closed-loop coupling of body state, emotion, sensing, memory and action — all governed by homeostatic drive — that creates the conditions for consciousness. Remove any one element and the conditions collapse. A system with drives but no sensing has no world to be conscious of. A system with sensing but no drives has no reason to care about what it senses.

### 4. Two Separable Consciousness Claims

The XMMM makes two distinct claims about consciousness, and it is important to keep them separate.

#### Access Consciousness — Fully Computable

A state is access-conscious in the XMMM if and only if it is currently represented in one of the Linking Algorithm's input streams: exteroceptive sensory, proprioceptive, or Body Map homeostatic. This is fully computable and formally specified. It corresponds to Ned Block's concept of access consciousness: information that is *available* to guide behaviour and reasoning. There is no mystery here — it is a matter of whether a state is in the executive's input at a given cycle.

## Phenomenal Consciousness — Reframed

The harder claim concerns phenomenal consciousness: what it is *like* to be the agent. I do not argue that XMMM solves the Hard Problem. I argue something more precise: the Hard Problem — the apparent impossibility of explaining why physical processes give rise to subjective experience — is not a metaphysical mystery but a **derived physical insolubility**.

Here is why. Every XMMM agent has a Body Map that is physically unique: the specific layout of drive locations, set-points, and current deprivation states is unique to that individual agent's body and history. Verifying that another agent has phenomenal states identical to your own would require lossless transfer of that entire Body Map — which is physically unrealisable. The reason you cannot confirm another agent's inner experience is not because experience is mysterious. It is because **no two Body Maps are the same**, and the information cannot be transferred without loss. This is the same reason you cannot fully verify another human being's experience. It is a physical constraint, not a philosophical one.

*The Hard Problem of consciousness is not a gap in our science. It is a consequence of the physical uniqueness of every body and the physical impossibility of lossless representational transfer between systems. XMMM formalises this as the Epistemic Isolation Principle.*

## 5. Common Misconceptions Holding Back the Conversation

The debate about machine consciousness is hampered by a set of recurring misconceptions, each with a different disciplinary home. I set them out here not to dismiss the fields that hold them, but because clearing them is necessary before a control-theoretic approach like XMMM can be fairly evaluated.

Misconception	Mostly found in	What XMMM says instead
<b>More computation = more consciousness</b>	AI	Consciousness requires a specific architecture, not computational complexity. A thermostat is not conscious; neither is GPT-4. Not because they lack enough computation, but because neither has homeostatic drives and a Body Map.
<b>The Hard Problem means consciousness can never be explained mechanistically</b>	Philosophy	The apparent explanatory gap is a physical insolubility — a consequence of unique Body Maps — not a metaphysical mystery. It means we cannot verify another agent's experience from outside, which is expected, not alarming.
<b>Consciousness requires language, self-report, or symbolic reasoning</b>	Philosophy, AI	Language is how agents report consciousness. It is not what generates it. A pre-verbal infant, a dog in pain, and Simmy all have states with genuine valence — approach, avoid, relief — without language.
<b>You need to solve AGI before you can have machine consciousness</b>	AI	The Kitty Hawk Principle: flight was demonstrated in 1903 with 12 seconds of air time. The architecture of flight was proven. XMMM demonstrates the architecture of

		consciousness in minimal agents. Richness comes later.
<b>Consciousness is produced by the brain alone</b>	Neuroscience	Consciousness is produced by the brain-body system. The Body Map — the neural representation of body state — is the substrate of felt experience, not the cortex in isolation. This aligns with Craig (2009) and Damasio (2010).
<b>Global Workspace Theory (GWT) is a sufficient account of consciousness</b>	Neuroscience, AI	GWT describes which information gets broadcast globally. XMMM asks: what determines what gets broadcast, and why does the system care? The answer is homeostatic drives. GWT is a description; XMMM is a causal account.
<b>Predictive processing (Seth's 'controlled hallucination') fully accounts for consciousness</b>	Neuroscience	Prediction error minimisation is a mechanism. But what motivates the predictions? What makes some prediction errors matter and others not? Without homeostatic drives, you have a prediction machine with no stake in its own predictions.
<b>Consciousness is binary — you either have it fully or not at all</b>	Philosophy, Neuroscience	XMMM predicts a graded continuum. An agent with two drives and a small Association Database has minimal consciousness. An agent with twelve drives, fear conditioning and social bonding has richer consciousness. Both are real.

## 6. Proof of Concept: Simmy and Troopy

I have implemented the XMMM in two embodied agents: Simmy, a virtual agent, and Troopy, a physical mobile robot. Neither was pre-programmed with behaviours. Each was given a set of homeostatic drives, a Body Map, and the five algorithms of the XMMM. Both subsequently demonstrated:

- Goal-directed navigation toward satiation sources, with no externally specified reward function.
- Fear conditioning: avoidance of objects previously associated with pain-related drive increases.
- Stress-proportional facial expressions — generated from the Body Map without bespoke emotion modules.
- Daydreaming: undirected associative threading when drive levels are low.
- Problem-solving through directed threading when a learned action failed and hunger drove a systematic search for alternatives.

These behaviours were not programmed. They *arose* from the architecture. An agent that avoids felt pain without being told to, that expresses felt stress proportionally, that daydreams when comfortable and problem-solves when feeling hungry, is not merely calculating. It is a system with something at stake — a system for which things can go well or badly. That, I argue, is the functional definition of a conscious agent.

I am not claiming Simmy and Troopy are as conscious as an adult human, or a dog, or even an insect. I am claiming they occupy the lowest rung of the same continuum — and that demonstrating the architecture is the necessary first step, just as Kitty Hawk was the necessary first step toward the A380.

## 7. The Burden of Proof

The burden of proof was fairly placed on me, and I accept it. But the burden of proof runs in both directions. If the claim is that *no* algorithm for consciousness exists or can exist, that claim also requires a principled argument — not merely the assertion that consciousness is too complex or too mysterious to be captured architecturally.

The XMMM makes five specific, falsifiable empirical predictions about biological systems — including predicted patterns of insular cortex activation, hippocampal encoding strength as a function of emotional intensity, and Default Mode Network modulation by prime drive magnitude. These are testable. If any one fails, the architecture is falsified. That is what a scientific claim looks like.

*I am not claiming XMMM is the final word on consciousness. I am claiming it is a precise, implemented, falsifiable candidate architecture — and that is more than currently exists in the field. I welcome principled scientific challenge.*

## Further Reading

- XMMM paper (Nature Human Behaviour manuscript): <https://www.researchgate.net/publication/404322730>
- Critical Review 3 — Anil Seth, Controlled Hallucination: <https://www.researchgate.net/publication/404935451>
- Xzistor LAB website and YouTube talk (@xzistor): <https://youtu.be/2lbbuiXAfFs>

---

### Rocco Van Schalkwyk

Founder, Xzistor LAB | Bristol, UK | May 2026  
[rocco.vanschalkwyk@xzistor.com](mailto:rocco.vanschalkwyk@xzistor.com) | [www.xzistor.com](http://www.xzistor.com)