



The Synthetic Nervous System


Architecting Distributed Edge AI, Autonomous Robotics,
and Ephemeral Intelligence on a \$15 Microcontroller.



THE VESSEL AND ITS CONSTRAINTS




THE BRAIN: 8 MB SPI FLASH LIMITS



THE EDGE AI PIPELINE: SHRINKING THE MIND

1. CLOUD TRAINING

35 specification:
20 KB of 260 images



Heavy compute derives mathematical weights from thousands of images.


2. QUANTIZATION



Shrinking heavy 32-bit floating-point data into dense 8-bit integers (INT8). Reduces footprint to 100 KB - 1 MB.

3. EMBEDDING

The tiny TFLite model is compiled directly



The tiny TFLite model is compiled directly into the 8 MB Flash memory alongside firmware.


4. REAL-TIME INFERENCE



TFLite micro-interpreter executes in 8 MB PSRAM using hardware vector instructions.

HIERARCHICAL COGNITION


Bypassing flash memory limits by storing 50+ specialized 1 MB models on a 32 GB SD card.



Execution Strategy: The board incurs the heavy compute and read-time cost of a specialist model only when a high-confidence primary target justifies the energy expenditure.

TEMPORAL PERSISTENCE: HACKING HARDWARE WITH PHYSICS

Reading a 1 MB model from SD via SPI takes up to 600ms, which would normally crash the frame rate. We solve this bottleneck by assuming the physical persistence of the target.



**Frame 0: Trigger
(Cat Detected)**

**The Loading Window:
300ms SD Card Read**

System intentionally ignores incoming frames, using CPU purely to load the specialist model into PSRAM.


High-Speed Analysis

Model runs flawlessly in fast memory. A physical cat cannot teleport out of frame in 3 seconds, allowing top-speed inference without re-reading the SD card.

SPATIAL EFFICIENCY VIA DYNAMIC ROI CROPPING


Instead of processing dead space, the base model extracts only the Region of Interest (ROI) via a fast PSRAM pixel copy.

FULL FRAME PROCESSING




FULL 640X480 FRAME
921 KB RAM FOOTPRINT
PROCESSING SPEED: 5 FPS

CROPPED ROI PROCESSING




CROPPED 128X128 BOUNDING BOX
49 KB RAM FOOTPRINT
18X MEMORY REDUCTION
PROCESSING SPEED: 20+ FPS

MUULTIMODAL FUSION: ANCHORING PERCEPTION





THE ACTIVE LEARNING LOOP (AUTOPOIESIS)

The architecture transforms from a static sensor into a self-improving surveillance network.




THE ARCHITECTURE OF NETWORKED INTELLIGENCE

	Local Inference (Static Offline)	Full Video Streaming (Traditional)	Split Vector Streaming (Modern Paradigm)
 Wi-Fi Bandwidth	Minimal (Alerts only)	Heavy (1-5 Mbps)	Ultra-Low (~4 KB/s)
 Edge Compute Load	Critically High	Low	Medium (Feature backbone only)
 Server Compute Load	Zero	High (Decoding overhead)	Low (Pure vector math)
 Data Privacy	High	Low (Raw video exposed)	High (Geometric data only)


SPLIT INFERENCE: THE GEOMETRY OF TRUTH

The edge device strips away lighting and background noise, acting solely as a feature extractor. It outputs pure mathematics.




EPHEMERAL INTELLIGENCE AND NEUROPLASTICITY

Narrative insight: The Server evaluates vector streams and dynamically rewrites the physical mindset of the Edge in real-time.



EMBODIMENT: THE SPINE VS. THE BRAIN

Resolving the Bipedal Payload Problem: Heavy onboard compute destroys balance.
We offload heavy matrix math to keep the torso lightweight.




SKIN IN THE GAME: THE Q7 SENTINEL

Transforming a commercial robot vacuum into an intelligent, roaming sensory platform.

REROOTING ARCHITECTURE:

- UART bootloader interception flashes open-source Valetudo firmware.
- Liberates device from cloud constraints.
- Native LiDAR handles autonomous navigation.
- Server identifies visual anomalies and commands Q7 to intercept the target.



ORCHESTRATING THE HEARTBEAT: MANAGING PHYSICAL CONSTRAINTS

Resolving the challenge of maintaining a reliable core reflex system under heavy concurrent processing loads.

THE HARDWARE REALITY

STADTOFA: RETNS PBB, 1 20.21

TECHNICAL CRASH
04% : 00



RAM EXHAUSTION

Silent crashes when state-tracking vanishes.



CPU CONTEXT SWITCHING

Logic stutter under heavy polling.



VOLTAGE SAG

Peak current draw destabilizing logic gates.

AUTASCTRAHDOEA
04% : 00

THE ARCHITECTURE FIX


KSTGUETHAK ORLL KEAS, 1 10.32

THE SUPER-LOOP



Single-threaded execution crashes under concurrent logic loads.

RTOS PRIORITY QUEUE



Real-Time Operating System decoupling ensures the core reflex heartbeat is never starved of CPU cycles by background tasks.



INTEGRAL SYNTHESIS: ARCHITECTING AN EVOLUTION

From fragmented data to integrated perception. We are establishing the blueprint for a synthetic nervous system.

SUBJECTIVE (INTERIOR)

FLUID CONSCIOUSNESS
Dynamic mindset rewriting via
Ephemeral Intelligence and
Micro-models.

INDIVIDUAL





EFFICIENCY & PHYSICS
Overcoming the Energy-
Bandwidth Paradox via Split
Vector Streaming.

COLLECTIVE

THE SYMBIOTIC DANCE
Bidirectional co-creation
of understanding between
the Server and the Edge.

061 : 88



AUTOPOIESIS
The self-improving physical
infrastructure of the active
learning network.

061 : 88

OBJECTIVE (EXTERIOR)

The system learns to inhabit its world with increasing grace.