




AWARE WHILE THE SYSTEM SLEEPS

SUBCONSCIOUS AI ON GPX10 PRO



An architecture for continuous multimodal awareness and significance driven compute

Continuous awareness should not require the entire system to remain fully alert.

Awareness continues during sleep


When a person falls asleep, conscious attention recedes, but sensory awareness does not disappear. The brain continues to process selected sensory signals from the surrounding environment. A loud noise or the sensation of something crawling on the skin can wake the body and prepare it to respond.

This does not mean that the sleeping brain sustains the same state of alertness used for conscious thought and deliberate action. Instead, it operates in a reduced state, continuously monitoring for potentially important signals and engaging higher levels of response when a signal becomes important.

GPX10 Pro applies the principle in silicon. A programmable Always-On (AON) intelligence layer can continuously acquire sensor data, extract features, combine context across multiple streams and perform neural network inference while the Arm subsystem is in sleep state and system resources are inactive. When the detected context requires a higher-level processing, the AON domain can generate an interrupt or wake-up event to Arm subsystem to wake and take control.

The comparison is architectural rather than biological. GPX10 Pro does not attempt to reproduce human cognition. It borrows one useful idea from the brain: awareness and full alertness do not need to consume the same resources.

Subconscious AI brings this operating principle to an edge device. Stay aware. Decide what matters. Become fully active only when needed.



THE SYSTEM PROBLEM

Conventional systems wake before they understand

Many battery-powered products spend most of their lives waiting. Their sensors continue to produce audio, motion, physiological or environmental data, while only a small fraction of these signals represent events that require complex inference, a user alert or a radio transmission.

In a conventional low-power architecture, the processor can remain in a sleep state while a sensor or peripheral monitors for predefined events. When an interrupt occurs, the processor must wake to read and interpret the available data. This introduces energy overhead from waking the CPU, restoring system state, accessing memory, moving data, and executing software—even when the detected event ultimately proves insignificant.

A more efficient architecture keeps a small, programmable intelligence layer active in the low-power domain. This layer can continuously process selected sensor data, determine whether an event is meaningful, and wake the main processor only when higher-level processing is required.

The architectural question

Can an edge device remain aware of its environment without keeping the full processor, memory subsystem and peripherals active? The answer requires more than an efficient neural accelerator. It requires a programmable intelligence layer within an always-on, ultra-low-power domain that can continuously monitor the environment, detect meaningful events, and selectively wake the higher-performance system when needed..

System activity should scale with the significance of the event rather than the arrival of sensor data.

Subconscious AI on GPX10 Pro

GPX10 Pro separates continuous awareness from full system activity through two principal power domains. The Always-on (AON) domain contains the sensing and programmable AI resources needed for continuous monitoring/interpreting the environment. The main compute domain includes the Arm processor, five additional MX8 AI cores, memory and system interfaces that can be activated when a workload requires higher compute capacity or broader system functionality.


Subconscious AI domain

- ✓ Programmable AON AI cores (Five MX8 AI cores)
- ✓ Low-power Sensor DMA
- ✓ Sensor-data buffers / 10 sensor-data tanks
- ✓ Ultra-low-power sensor ADC
- ✓ I2S audio pipeline
- ✓ Sensor fusion orchestration
- ✓ 256-point FFT and signal preprocessing
- ✓ Sensor data conditioning & feature extraction
- ✓ Event / threshold detection
- ✓ Low-power clocking
- ✓ Wake-up / interrupt generation
- ✓ 120 KB of programmable memory access.
- ✓ Scratchpad and feedback loop for NN staged inference

Full system resources

- ✓ Arm Cortex M4F processor
- ✓ 5x additional MX8 AI cores
- ✓ 2 MB SRAM / Program/ Model Memory
- ✓ Flash / Program Memory Interface
- ✓ Advanced Sensor & Data Processing
- ✓ Multi-Sensor Data Fusion
- ✓ Unified system memory
- ✓ Additional interfaces and peripherals
- ✓ Connectivity/ communication Interfaces
- ✓ Data Logging / Result Management
- ✓ Interrupt & Event Handling
- ✓ Application-directed voltage and frequency scaling

Before the device enters its low-power operating state, the system loads and configures application prepares the models, weights and event-detection logic needed by the AON domain. The AON subsystem can then continuously acquire, process and interpret sensor data while the larger main compute subsystem remains unavailable. Routine conditions stay local. Significant conditions generate a wake trigger.



MULTIMODAL AWARENESS

Multiple signals become context

A sleeping person does not respond to every sound or touch in isolation. The brain combines sensory information with context before deciding whether to remain asleep or become alert. Subconscious AI uses the same high-level principle: a single sensor can detect activity, while multiple sensors can help the system understand what that activity means.

GPX10 Pro can organize up to ten buffered sensor streams inside the AON path, subject to the selected interface configuration. Sensor data may be acquired and buffered independently, then processed down the appropriate signal-processing chain for fusion. AON resources may be used to perform filtering, filter-bank analysis, and 256-point FFTs to generate compact representations of the sensor waveforms. The extracted features from the sensor data may then be organized and aligned within a common observation window for temporal correlation across sensors. Finally, the processed sensor features may be combined and provided to the neural-network inference engine, allowing the AON domain to make decisions based on contextual information without requiring the main compute subsystem to be woken for every sensor event.

Resource scaling and performance scaling

Two different power-management decisions exist in this flow. The resource- scaling decision consists of determining which power domains, compute resources, and interfaces should be powered up to address the workload. Dynamic voltage and frequency scaling (DVFS) then takes over, deciding what voltage to apply and at what frequency to run the resources that have been powered up.


Decoupling these 2 mechanisms clarifies the energy-management strategy: AON intelligence is responsible for deciding whether the workload requires higher levels of system activity, whereas the main compute subsystem + its power-management mechanisms are responsible for figuring out what operating point best suits the workload that needs to be executed.

APPLICATION JOURNEY ONE

Smart rings that adapt PPG in real time

Many smart rings run the PPG sensor at a baseline sampling rate or schedule that is sufficient for routine activity while protecting battery life. A workout can justify a higher sampling rate so the ring captures richer optical data from the beginning of the session. If workout classification is offloaded over BLE, the phone round trip and application scheduling can delay the command to increase PPG sampling. By the time the ring responds, part of the most valuable onset window may already be gone.

Subconscious AI keeps a compact workout detector active on-device using low-power IMU motion patterns. Because the detector runs continuously rather than in widely spaced duty-cycled windows, it can recognize the transition into exercise locally and immediately raise the PPG sampling rate. Reflex Compute can then activate any additional processing needed for workout classification or richer analysis. When the workout ends, the application returns the PPG and compute resources to their baseline operating state.



APPLICATION JOURNEY TWO


Respiratory alerts that escalate without waiting

Why the first layer must remain on

Consider a wearable designed to detect signs consistent with a possible asthma attack. In a severe event, avoidable inference delay can reduce the time available to alert the user or caregiver. A phone-dependent detection loop also relies on radio availability, operating-system scheduling and connectivity, while aggressive duty cycling creates blind intervals. The first detection layer therefore needs to run locally and continuously.

Sensitivity first then specificity

The Subconscious AI model is not intended to make the final clinical determination. Its job is to continuously monitor the relevant signals with high sensitivity, and wake Reflex Compute. The higher-capacity model can then process a richer set of features and compute to improve specificity before application logic issues an alert or takes another approved action. A sensitivity-biased first stage may trigger more often than strictly necessary; the system can still reduce average energy if the low-power model dominates monitored time and the cost of occasional escalations is lower than running the larger model continuously.



**The always-on layer protects the response window.
Reflex Compute supplies the higher-confidence decision.**

A platform built for the complete cycle

GPX10 Pro combines ten MX8 AI cores, a programmable Always-on (AON) domain, configurable sensor-data tanks, on-chip signal/data preprocessing, multimodal AI inference and an Arm-controlled system response. The architecture enables OEM to design for continuous, low-power awareness first, and then selectively activate additional compute and communication resources based on the significance of the detected event.

Remain aware without remaining fully awake.

Request a GPX10 Pro evaluation

Bring a representative sensor configuration, first-stage and confirmatory models, and representative event profiles. Ambient Scientific can help map the continuous workload into Subconscious AI domain, define the Reflex Compute response and measure detection latency, sampling transition, false escalations and energy across the complete application journey.

ambientscientific.ai | Contact Ambient Scientific

EDITORIAL NOTE: Before release, add the approved evaluation URL and replace journey benchmark placeholders with measured values, including workload, sensor mix, model size, sampling rate, operating point and test condition. Keep medical claims within approved clinical and regulatory evidence.

Draft sources GPX10 Pro Datasheet 2026 | Ambient Scientific A Unique Analog Digital Architecture for Low Power AI