



BONAI

SECURITY

Lumen  S

<https://lumenostechnologies.com>

Table of Contents

1. Executive Summary
2. Multi-Camera Management
3. Real-Time Object Detection
4. Pose Estimation & Fall Detection
5. ANPR – Automatic Number Plate Recognition
6. Cloud Intelligence
7. Custom Detection Zones
8. Motion Detection & Tracking
9. AI Control Modes
10. Analytics & Logging

Executive Summary

You get a motion alert. Your heart spikes. You pull up the feed... and it's a dog. A shadow. A cloud. Or worse, your camera caught the shape of a person, but missed the knife in their hand because its AI only knows two classes: "Human" and "Car."

That's not security. That's a crying-wolf machine that desensitizes you to the one moment that actually matters.

Here's the hard truth: Standard on-camera AI is a dumb binary filter. It asks one question: "Is it a person or a vehicle?" If the answer is yes, it screams. It has zero context, zero object granularity, and zero understanding of intent.


It sees a fallen elderly person and classifies it the same as a gardener watering the lawn. It sees a license plate and just calls it a "car."

You are paying for a motion detector with a fancy label, not for intelligence.

Enter LumenOS BonAI. We don't just detect motion; we infer intent.

We ripped out the single-purpose logic and replaced it with a hybrid multi-modal inference engine that runs right in your browser, fusing on-device speed with cloud-grade depth.

Multi-Camera Management



Add cameras via HTTP/MJPEG, RTSP (via proxy), RTMP, or local video files.

Supports multiple grid layouts (1x1, 1x2, 2x2, 3x3, PiP).

Each camera gets its own live feed, controls (play, stop, snapshot), and processing pipeline.

Real-Time Object Detection

- We don't just see a “Person” – we see the threat vector.
- Our object recognition pipeline tracks 80 distinct object classes and weights each one dynamically. A dog is a low-threat ping. A backpack is medium. But a knife or scissors? That triggers an immediate high-threat critical alert before the subject even crosses your perimeter.



TOTAL 432 HIGH 581 MED 1819 LOW 240 FALLS 0										ZONES HIT 0
22:31:19	12566041-UHD_3840_21	person 76%	person 68%	person 63%	person 61%	person 55%	car 54%	car 53%	person 53%	
22:31:19	15396218_1920_1080_2	car 91%	person 66%	traffic light 62%						
22:31:19	12566041-UHD_3840_21	person 76%	person 68%	person 63%	person 61%	person 55%	car 54%	car 53%	person 53%	
22:31:18	15396218_1920_1080_2	car 91%	person 66%	traffic light 62%						
22:31:18	12566041-UHD_3840_21	person 76%	person 68%	person 63%	person 61%	person 55%	car 54%	car 53%	person 53%	
22:31:18	15396218_1920_1080_2	car 91%	person 66%	traffic light 62%						
22:31:18	12566041-UHD_3840_21	person 76%	person 68%	person 63%	person 61%	person 55%	car 54%	car 53%	person 53%	

Each detection is assigned a confidence score and a threat level (low/medium/high) based on the object type and confidence.

Pose Estimation & Fall Detection

We read the room, literally.


While legacy systems ignore posture, MoveNet Lightning runs real-time torso-orientation analysis.

We don't just see a person standing; we flag the torso-width-to-height ratio that indicates a collapse.

A Fall Event is logged and it immediately overrides every other signal and pushes a priority notification, because seconds matter when someone is down.



ANPR – Automatic Number Plate Recognition



We turn vehicles into actionable data.

Every car, truck, or bus triggers our heuristic ANPR layer. We don't just say "a car arrived." We scrape the plate, run it through a live alphanumeric OCR filter with confidence scoring, and log the unique identifier. You know who arrived, not just that something arrived.

Cloud Intelligence

Sends a JPEG frame plus detection/zone context to Google Gemini for scene understanding.

Returns:
Scene description
Threat list
Risk score (0-100)
Recommended action.


These are fused with on-device threat weights to produce a unified Intelligence Score.

Audio alerts (synthetic tones) for high-risk events.

Custom Detection Zones

Object-aware zones – alerts only when a detected object enters a defined polygon/rectangle.

1. Draw rectangular zones on a canvas overlay.
2. Zones can be assigned per camera or globally.
3. When an object enters a zone, an alert is triggered and logged.
4. Zones are saved in Storage.



With custom detection zones, you decide where context matters. A person in the driveway is routine; a person in the restricted perimeter zone at 2 AM with a high-threat object is a showstopper. No more blind spots, just intelligent, geofenced awareness.

Motion Detection & Tracking

Every frame is compared to the previous one using a pixel-threshold matrix (adjustable 1-100). We only care about changes that exceed your sensitivity floor, meaning a leaf rustle doesn't even register, but a walking human creates a dense cluster of differential pixels.

Those clusters are then grouped by proximity into tracked objects with centroid coordinates, size, and a threat label (low/medium/high based on pixel mass).

Each camera maintains a rolling 180-second motion history (motionData), a live pulse graph that shows you intensity over time, not just a binary ON/OFF switch.

This feeds directly into your dashboard's Motion Activity chart, giving you a visual heartbeat of your entire property across the last minute.



Heat Map Engine – The Spatial Memory

This is where we turn motion into institutional knowledge.

We overlay a dynamic grid (adjustable cell size) over each camera's field of view. Every time motion is detected in a cell, that cell's heat value accumulates. But it doesn't just grow forever, it decays exponentially by 5% per frame, so recent activity glows hot, while stale activity fades into the background.

This creates a living, breathing thermal signature of your environment. You aren't looking at a static map; you're looking at a real-time attention map that shows you exactly where the energy of your space is focused.

The Bottom Line

Standard motion detection gives you a fragment.

LumenOS gives you a map, a clock, and a memory.

- Motion Activity tells you when things happen.
- Heat Mapping tells you where things keep happening.
- Combined, they tell you what's about to happen because patterns always precede events.

AI Control Modes

AI Control Modes

Scan Interval: How often to run object detection (2–30 seconds).

Continuous Mode: Runs AI every N frames (1–10) for near-real-time detection.

Trigger Modes: Auto, On Motion, or Manual.

Global toggles for Vision AI, Pose AI, and ANPR across all cameras.

DETECTION & AI CONTROLS

SENS 6 · V-AI 4 CAM · ∞ 4 · POSE 4 · ANPR 4

MOTION DETECTION

SENSITIVITY

PIXEL THRESHOLD

OVERLAYS

BOXES

ZONES

LABELS

GRID

POSE

RESET ALL DETECTION

VISION AI SETTINGS

SCAN INTERVAL

5S

CONTINUOUS FRAMES BETWEEN SCANS

3

MODE

SECURITY

DETAILED

TRIGGER

AUTO

ON MOTION

MANUAL

POSE AI · FALL DETECTION (MOVENET)

POSE SCAN INTERVAL

5S

FALL CONFIRM FRAMES

2

MIN KEYPOINT CONFIDENCE

0.30

MoveNet Lightning – torso-orientation based fall detection.

GEMINI INTELLIGENCE SETTINGS

GEMINI COOLDOWN

8S

Min seconds between Gemini calls per camera. Audio alerts fire at HIGH (50+) and CRITICAL (75+) risk scores.

ANPR · PLATE OCR (TESSERACT.JS, ON-DEVICE)

ANPR SCAN INTERVAL

4S

PLATE BAND HEIGHT

30%

MIN OCR CONFIDENCE

55%

Crops the lower-centre band of each detected vehicle box and runs OCR on it. Heuristic region-of-interest – no dedicated plate detector model.

Analytics & Logging

Dashboard shows live status, AI activity log (with threat levels), Gemini feed, ANPR log.

System alerts and event timeline.

Exportable reports (JSON, CSV, TXT) including camera states, detection logs, Gemini analyses, and ANPR reads.





Thank you!

<https://lumenostechnologies.com>