



Spatial intelligence for decisions that depend on where.

Different worlds. Same question. Where? What if spatial decisions took seconds, not hours?

A supply chain on Earth and a landing site on the Moon look like different worlds. They begin with the same question: where is the risk, where should we act, where is it safe, where could we stay. Anassa.ai is an independent research lab studying how autonomous AI agents can turn that question into decisions, at machine speed and under human governance. The framework is called ANASSA: Autonomous Neural Agents for Spatial Systems Architecture.

THE PROBLEM

Spatial data has never been more abundant, or more fragmented. Geospatial platforms, satellite feeds, sensors, cloud databases and legacy systems each operate alone, and people become the connective tissue between them: a fragmentation tax paid in analyst hours while the question that matters waits.

THE FRAMEWORK

One orchestrator coordinates more than 47 specialized agents, organised into eleven components across four layers, through a continuous six-step cognitive loop: Sense, Reason, Plan, Decide, Execute, Learn. Mandatory human review gates sit at Plan and Decide. Every spatial and numerical output is validated against authoritative ground truth before it can influence a decision. The layer sits above existing geospatial platforms, data and compute and connects through MCP, A2A, Geospatial SDKs and APIs, IoT connectors, legacy adapters, and others.

EXAMPLE USE CASES

Where should we build, protect, or intervene?

[Cross-source data reconciliation](#)

Measure: Pipelines that complete without a manual restart; records checked against the authoritative source.

Where is disruption emerging?

[Supply chain risk](#)

Measure: Minutes from a hazard update to a refreshed risk map; analyst hours per assessment.

Where is it safe to move and act?

[Autonomous vehicle safety](#)

Measure: Simulation coverage gained per unit of compute.

Where could humanity establish a lasting presence?

[Lunar site selection](#)

Measure: Specialist review time; every numerical claim traced to a primary source.

THE STORY BEHIND THE LAB

"I spent years as a technical program manager, the human glue between tools that don't talk. When agentic systems arrived I argued the role wasn't disappearing, it was moving upstream: to writing intent, setting guardrails and making the calls that still need a human. Spatial decisions carry the same fragmentation tax at far higher stakes. Anassa.ai is that observation turned into a research programme. I've been the glue; this is what I think replaces it."
— Constantinos Papantoniou, Ph.D., Founder and Execution Architect

PUBLICATIONS AND NEXT STEP

White paper (June 2026), high- and low-level architecture designs (v1.0) at anassa.ai. Two applied research papers with Brian Hilton, Ph.D.: Supply Chain Deployment (in review) and Earth and Lunar Applications (in progress). Next: implementation-level validation with a partner on a scoped, measured pilot. Enterprise pilots, research collaboration and academic partnerships welcome.