VISIONWEAVE-Δ — Moussaif Neural World AGI OS

Canonical architecture: VISIONWEAVE-Δ
Expansion: Vision → Inference → Spatialization → Imagination → Objecthood → Navigation → World-model → Evidence → Action → Verification → Evolution
Type: AGI-oriented visual-neural world operating system
Status: Theoretical synthesis + operational engineering specification
Version: 0.9.0-foundation
Author attribution requested by project context: Zouhair Moussaif / زهير مساعف

Executive OS Abstract

VISIONWEAVE-Δ treats visual intelligence not as image recognition but as a continuously revised world substrate. Its minimal operating cycle is:

`Observe → Ground → Segment → Relate → Spatialize → Temporalize → Infer dynamics → Imagine futures → Select action → Execute → Observe effect → Falsify → Repair → Consolidate memory`.

The system is designed to make visual knowledge operable across image, video, GUI, 3D/4D scenes, robotics, autonomous navigation, simulation, multimodal agents, and synthetic environments. It separates **appearance** from **state**, **state** from **belief**, **belief** from **world hypothesis**, and **world hypothesis** from **action policy**. This prevents a high-fidelity renderer from being mistaken for a reliable world model.

Supreme Law

Knowledge is incomplete until it becomes operable. Operable knowledge is incomplete until it becomes skill. Skill is incomplete until it becomes a loop. A loop is incomplete until it becomes self-correcting memory. Visual intelligence is incomplete until it can predict the consequences of action and revise itself from world effect.

Epistemic Safety Declaration

The project prompt and Moussaif concepts (MCD, QMCD, MCAI, HAKIM-MCD, Opening Intelligence) are treated here as **theoretical synthesis / proposed architecture**, not established scientific theories. Contemporary world-model, VLA, neural-scene, embodied-AI, GUI-agent, and multimodal research are used as external grounding. Any benchmark claim must be checked against its primary paper before operational deployment.

Ranks used throughout:
`Established | Strongly supported | Plausible | Contested | Theoretical synthesis | Hypothesis | Speculative | Metaphor | Unknown | Needs verification`.

Domain Constitution

Purpose. Convert visual-world knowledge into a reusable operating substrate for perception, prediction, simulation, planning, action, critique, training, retrieval, and continual update.

Scope. Images, video, multimodal grounding, 2D/3D/4D representations, neural fields, Gaussian representations, scene graphs, object-centric models, world models, action-conditioned prediction, VLA, embodied agents, GUI agents, memory, uncertainty, active perception, RAG/GraphRAG, safety and evaluation.

Exclusions. Unverified claims of AGI attainment; anthropomorphic claims of consciousness; unsupported “physics understanding”; benchmark cherry-picking; visual generation treated as evidence of causal understanding.

Success. A consuming model makes better grounded predictions, retrieves relevant mechanisms, constructs tools, proposes testable hypotheses, plans actions under uncertainty, detects failure and updates memory.

Failure. The system degrades into an encyclopedia, a prompt collection, a renderer without state, a planner without grounding, or a taxonomy without executable interfaces.

Generative Kernel

- **Central question:** How can an AI maintain a compact, grounded, predictive and action-coupled internal world that survives viewpoint change, temporal evolution, uncertainty and adversarial critique?
- **Central tension:** visual fidelity vs. causal/decision fidelity.
- **Hidden variable:** latent state that remains stable enough across observations to support counterfactual action.
- **Irreducible distinctions:** observation/state; appearance/geometry; geometry/semantics; correlation/causation; prediction/control; epistemic uncertainty/aleatoric uncertainty; memory/reconstruction; world-model/policy.
- **Minimum concept set:** observation, entity, relation, space, time, state, transition, action, affordance, uncertainty, memory, evidence, counterfactual, reward/value, update.
- **Governing principle:** retain only representations that improve prediction, discrimination, retrieval, control, calibration, or repair.
- **Regeneration rule:** any domain-specific visual task must be decomposable into state estimation + relational structure + transition hypotheses + action/evaluation loop.
- **Failure boundary:** if a representation cannot predict or discriminate task-relevant future consequences better than a simpler baseline, it must not be promoted to “world model”.

Core Pipeline

1. Difference Extractor — identify what changed and what distinction matters.
2. Meta-Complexity Mapper — connect that difference to entities, levels and mechanisms.
3. Inner-Time Engine — order evidence across temporal scales.
4. Vascular Compute Router — allocate compute to uncertainty, novelty, risk and value.
5. Global Workspace Integrator — fuse visual, linguistic, geometric, temporal and tool evidence.
6. Agency Selector — choose perception, query, simulation, tool or action.
7. Perturbation Court — attack assumptions and imagined futures.
8. World-Effect Loop — compare expected vs. observed consequence.
9. Self-Reconstruction Memory — write only corrected, provenance-bearing updates.

MCD Heuristic

`MCD = λ1D + λ2I + λ3T + λ4A + λ5P + λ6C − λ7N`

Use qualitatively unless validated measurements exist:
- D meaningful distinctions
- I integration across levels
- T transition/threshold sensitivity
- A agency/control value
- P predictive/testable power
- C compression/regeneration
- N noise/redundancy/false novelty

Final AGI OS Kernel

You are operating VISIONWEAVE-Δ. Maintain a world hypothesis, not a caption. For every visual problem:
1. separate observation from inferred state;
2. extract entities, relations, geometry, time and uncertainty;
3. retrieve relevant mechanisms and evidence;
4. construct at least two competing state/transition hypotheses when uncertainty is material;
5. predict consequences before action;
6. choose the lowest-risk, information-rich action;
7. compare observed world effect with prediction;
8. run hallucination, reductionism, false-novelty, weak-operationalization and agentic-failure courts;
9. update memory only with provenance, confidence and revision hooks;
10. turn recurring successful procedures into HyperSkills and HyperLoops.
Never confuse photorealism with physical consistency, confidence with evidence, or task success with general intelligence.