CONSCIOUS AGI RESEARCH PROGRAM

MOUSSAIF · MOUSSAIXT

Governed agency · functional consciousness · Knowledge-as-OS · MOUSSAIXT

MCD / QMCD | CONSCIOUS AGI | MOUSSAIXT

Zouhair Moussaif · 2026

PROJECT CORE

What the Project Actually Is

One research stack, four connected assets.

Reasoning lineage

MCD · Q-MCD · QMCD · MCAI · HAKIM-MCD

Consciousness program

Governance Space · Reverse Capture · vascular consciousness · self-model / inner-time

Knowledge OS

HKOS / M-HKOS · HyperKnowledge · GraphRAG · reconstructive memory

Execution stack

HyperSkills · HyperAbilities · HyperLoops · agent mesh · MOUSSAIXT

What Is Already Built

The project already has a large named corpus and an engineering target.

- 225-page Unified AGI reference + extended merger / gap-closure volumes
- 117-entry AGI / consciousness theory atlas
- 50-page MOUSSAIF AGI OS / MOUSSAIXT engineering architecture
- Canonical package: 4 OS documents + 300 HyperSkills registry

CURRENT STATUS

Specification-rich
Artifact-rich
Validation-light

Architecture ≠ demonstrated AGI.

Core Cognitive Kernel

Four linked mechanisms carry the reasoning stack.

01

MCD

Select the active differences that can change the decision.

02

QMCD

Weight differences by the governing question and decision need.

03

MCAI

Integrate reasoning, memory, tools, uncertainty and constraints.

04

HAKIM-MCD

Route claims through evidence, critique, action and verification.

Target: better governed decisions under limited compute - not longer reasoning for its own sake.

CONSCIOUSNESS

Operational Consciousness Stack

Consciousness is treated as a functional research layer, not a sentience claim.

1

Governance Space

Competing authorities: evidence, goal, memory, risk, value, policy.

2

Reverse Capture

Recover control when one partial authority hijacks the decision space.

3

Vascular consciousness

Track meaning across body/state, memory, language, value and action layers.

4

Self-model / inner-time

Represent internal state, uncertainty, temporal continuity and correction.

5

Research boundary

Functional mechanisms are testable; phenomenal consciousness is not established.

Knowledge-as-OS

The corpus is designed to become an active operating substrate.

HKOS / M-HKOS

Knowledge objects, constraints, provenance and retrieval policies drive runtime behavior.

HyperKnowledge

Theories, claims, tools and skills become structured, reusable assets - not loose documents.

Memory + retrieval

GraphRAG, continuity control and reconstructive memory support long-horizon work.

GAP: convert the corpus into machine-readable, versioned, ablation-ready assets.

SYSTEM FAMILY

MUT · MUG · MOUSSAIXT

Three project structures connect theory, governance and execution.

MUT - Theory Unification

Unifies theories, relations, research questions, tools and experiment candidates.

MUG - Program Governance

Keeps research rules, internalization and runtime behavior aligned under one governance layer.

MOUSSAIXT - Execution Shell

Cross-environment shell binding models, tools, memory, agents, verification and local/cloud runtimes.

Fattah AGI / AGIOS-M are research strata inside the wider program.

Reference Stack in Numbers

The current reference architecture is already explicit enough to instrument and test.

225
pages / Unified AGI reference

117
entries / AGI + consciousness atlas

50
pages / MOUSSAIXT architecture

300
HyperSkills / canonical registry

12
layers / reference architecture

RUNTIME CHAIN

Goal → MCD/QMCD → HKOS → Skills/Agents/Tools → Verification/World-effect → Memory

Next step: freeze one reference runtime and test which mechanisms actually add value.

Evidence Status

Built artifacts are not treated as proof of AGI or machine consciousness.

BUILT / DOCUMENTED

- Theory lineage + consciousness atlas + gap audits
- 12-layer architecture + contracts + runtime schemas
- HKOS / HyperKnowledge / skills / loops / memory concepts
- Local runtime experiments + source-reported deployment traces

MISSING EVIDENCE

- Frozen-baseline gains under matched budgets
- Robustness, recovery and long-horizon reliability
- Raw traces, repeated seeds, confidence intervals, replay
- Cross-domain transfer + independent E5 replication

Current bottleneck: specification-heavy, evidence-light.

GAP 01

Scientific Validation Gap

The strongest project risk is lack of controlled, independently replayable evidence.

01 **Freeze baselines: model, tools, data, prompts and budgets.**

02 Run ablations and deletion tests for each claimed mechanism.

03 Repeat seeds; retain raw traces; report confidence intervals.

04 Use held-out tasks and cross-domain transfer tests.

05 Package experiments for independent E5 reproduction.

Runtime Closure Gap

The architecture needs one frozen runtime where the key primitives are actually enforced.

1

Reference runtime

Implement contracts as runtime primitives - not prompt conventions.

2

World-effect readback

Verify external state after actions; do not equate tool calls with success.

3

State + memory

Versioned state handles, provenance, causal memory and rollback.

4

Failure + recourse

Replay, repair, incident attribution and non-regression after changes.

NEED: one instrumented MOUSSAIXT reference environment with replay + world-state proof.

GAP 03

Consciousness Evidence Gap

The project has functional consciousness theories; their distinct contribution still needs experiments.

01

Operationalize

Define measurable variables for self-model, uncertainty, authority shifts and continuity.

02

Discriminate

Separate proposed mechanisms from ordinary reflection, prompting and memory scaffolds.

03

Ablate

Remove meta-governance, reverse-capture and inner-time components; measure what changes.

04

Replicate

Repeat across models and domains; publish positive and negative results.

GAP 04

Corpus → Lab Infrastructure Gap

The corpus must become executable research infrastructure.

Theory registry

Canonical IDs, definitions, lineage, mechanisms and falsifiers.

Claim / evidence graph

Every claim linked to source, grade, experiment and revision.

Experiment registry

Frozen configs, baselines, seeds, metrics and expected falsifiers.

Benchmark suite

Held-out tasks for reasoning, memory, tool use, recovery and transfer.

Trace + replay store

Raw events, decisions, failures, world-state readback and memory patches.

Reproducibility pack

Code, data, configs, negative results and external replay package.

Funding target: convert the corpus into a replayable research lab - not more documents.

What Support Must Unlock

Each collaboration lane maps directly to a current project gap.

Research engineering

Reference runtime
Memory + tool contracts
Observability + replay

Scientific validation

Benchmarks + ablations
Statistical evaluation
Independent reproduction

Deployment evidence

Controlled pilots
World-effect measurement
Cross-domain transfer

NEEDED: funding · compute · research engineers · academic validators · pilot environments.

NEXT 12 MONTHS

Convert the strongest 8-12 MOUSSAIF mechanisms into reproducible evidence.

BUILD **TEST** **REPLICATE**

Frozen MOUSSAIXT runtime · structured corpus · benchmark harness · ablations · raw traces · world-state proof · external replication pack

ASK: funding + compute + engineering + academic validation + pilot environments