MOUSSAIF AGI

INTEGRATED FRAMEWORK

MUT × MUG × MOUSSAIXT

Theory corpus → capability placement → governed runtime → evidence

MCD | M-HKOS | MUG | MOUSSAIXT

A 15-slide map of the complete research logic — not a list of names.

Zouhair Moussaif · 2026

The framework is a closed research-to-runtime loop

MUT explains and organizes. MUG governs where capability lives. MOUSSAIXT executes and measures it.



Goal → Question → Difference → Evidence → Route → Plan → Action → Readback → Court → Memory → Better Question

Research integrity: architecture is only the test vehicle; every claimed gain must return as evidence.

SOURCE ANCHORS · Conscious AGI Master Corpus pp.3–8 · MUT-AGI-02 · MUG Tri-Body

One corpus — six functional strata

The names remain separate because they answer different research questions and carry different evidence boundaries.

1 · Reasoning

Governing Question · Highest Difference · MCD · Q-MCD · QMCD

2 · Meta-cognition

MCAI · MICAI · MICAGI · MICAGFI · NeoAGI · WMICAGI

3 · Consciousness

Governance Space · Reverse Capture · Vascular / Inner-Time / Self-Model

4 · Knowledge OS

M-HKOS · HKOS · Context Genome · provenance · reconstructable memory

5 · Agency

HyperKnowledge · HyperSkill · HyperAbility · HyperLoop · HyperAgent

6 · Execution + Proof

MUT · MUG · MOUSSAIXT · courts · world-effect · evidence ladder

Integrity rule · Preserve lineage → assign function → attach source → define test → rank evidence.

SOURCE ANCHORS · Unified AGI · Engineering Architecture · Consciousness Atlas · MUG/MUT Extension

MUT-AGI-02 is the corpus unifier — not another layer of names

Its value is the typed relation between theories, failure signatures, metrics, tools and experiment routes.

REGISTRY SCALE

42 core theories
20 legacy constructs
5 layer theories
200-node unified registry
100 operational tools
19 HyperKnowledge bases

COGNITION

Q-MCD · MCD · MCAI · HAKIM

KNOWLEDGE

M-HKOS · retrieval · memory · provenance

AGENCY

skills · loops · agents · tools · routing

EVIDENCE

courts · world-effect · benchmarks · evolution

What MUT adds: one place to ask “what does this construct change, how does it fail, and what would falsify it?”

SOURCE ANCHORS · MUT-AGI-02 master reference · 42/20/5 · 200-node registry · 100 tools · 19 HyperKB

MUT turns theory into an inspectable causal path

The core sequence starts with a mission contract and ends only after world-state readback and memory update.



THEORY CONTRACT

Each construct must have:
- object
- metric
- failure signature
- nearest baseline
- ablation / negative case

EXECUTION CONTRACT

No “done” on text alone.
Action → postcondition → readback → independent verifier → receipt → state update.

LEARNING CONTRACT

Only verified state changes become learning candidates.
Failure is retained as data, not hidden as retry noise.

MUT’s strongest claim is methodological: a theory becomes valuable when it changes a measurable transition.

SOURCE ANCHORS · MUT-AGI-02 · MOUSSAIFXT Part II update map · evidence contract

The 42-core registry is best read as families of mechanisms

The framework preserves the full registry while letting experiments target small bundles instead of testing everything at once.

CONSTITUTION + GOAL

MASO · M-AGIOS-K · MCIL · MGCG

QUESTION + DIFFERENCE

Q-MCD · MCD · Computational Vascularity

EVIDENCE + COURTS

MELT · Gap-Filling Court · Destruction Court · Mechanism Trace

WORLD + KNOWLEDGE

World-State · Project-World Graph · M-HKOS · Provenance Retrieval

MEMORY + CONTEXT

Stratified Memory · Context Capsule · reconstruction / continuity

AGENCY + CONTROL

HyperSkill Graph · Sparse HyperRouting · Bounded HyperLoop · Tool Intelligence · Plan Graph

Research strategy: select 8–12 highest-information constructs for the first E4 bundle, then expand only when ablations show marginal value.

SOURCE ANCHORS · MUT-AGI-02 registry TH-01…TH-28 · Master Reference pp.84–85

MUG separates three bodies that must never be conflated

It governs the translation between runtime behavior, training/internalization and explanatory theory.

BODY A · RUNTIME KERNEL

Goal · state · authority · memory · tools · verification · recovery.

Must remain inspectable, reversible and rapidly changeable.

BODY B · TRAINING STACK

Behaviors considered for adapter / weight internalization only after transfer, holdouts, retention checks and rollback readiness.

BODY C · UNIFIED THEORY

MUT's hypotheses about why mechanisms work, where they fail and what evidence would change the model.

Theory never receives runtime authority by name alone.

MUG's central problem is capability placement: where should a behavior live for reliability, cost, reviewability and reversibility?

SOURCE ANCHORS · MUG Tri-Body Governance · UT Austin theory-runtime-training bridge · EX-02

Capability is not “stored” in one place — MUG chooses the vessel

Context, knowledge, skills, loops, adapters and weights have different reversibility, cost and failure modes.

CONTEXT

Ephemeral task state

Promotion evidence: No persistence required

HYPERKB

Mutable provenance-rich knowledge

Promotion evidence: Repeated retrieval value

HYPERSKILL

Reusable governed capability

Promotion evidence: Cross-task wins + contract + tests

HYPERLOOP

Bounded multi-step workflow

Promotion evidence: Budget / stop / recovery structure

ADAPTER

Behavior portable across runtime/model

Promotion evidence: Stable gain with lower context cost

WEIGHTS

Mature high-frequency behavior

Promotion evidence: Holdout gain + retention + rollback

Promotion law · Do not move a capability into a less reversible vessel unless causal gain survives matched-budget ablation and rollback remains explicit.

SOURCE ANCHORS · MUG Vessel Decision Matrix · selective internalization · MATI

Selective internalization is an evidence-gated lifecycle

Start external. Promote only when stable transfer is demonstrated. Revert when evidence fails.



PROMOTION INPUTS

Evidence rank · volatility · privacy class · transfer score · target vessel · checkpoint / rollback target

GUARDRAILS

External by default.
No volatile/private facts in weights.
Matched-budget ablation.
Explicit checkpoint identity.

RESEARCH PAYOFF

“Should we fine-tune?” becomes measurable: compare success, latency, tokens, transfer, editability, forgetting and recovery across vessels.

MUG is therefore a governance theory of where capability should live — not a blanket preference for training or retrieval.

SOURCE ANCHORS · MUG EX-02 · VesselDecision schema · UT Austin sponsored-study design

M-HKOS is the knowledge substrate that MUT reasons over and MUG places

The shift is from “documents to retrieve” toward typed knowledge objects with provenance, contracts, skills, loops and writeback.

KNOWLEDGE OBJECT

Trigger · scope · source · claim/evidence · mechanism · skill/tool relation · failure authority · provenance

OPERATING CONTRACTS

Ontology / Concept Graph
Claim–evidence ledger
Retrieval policy
Skill / loop links
Tests · version · rollback

WHY IT MATTERS

A large corpus is useful only if the system can reconstruct the task state, retrieve the right object, preserve contradictions, and turn verified outcomes into typed writeback.

MUT supplies the theory graph. MUG decides what remains in HyperKB vs becomes skill / loop / adapter / weight. MOUSSAIXT executes the contracts.

SOURCE ANCHORS · MHKOS / KB-as-OS · MATI selective internalization · Unified AGI

Consciousness enters as a testable governance program — not a sentience claim

The corpus asks whether self-state, uncertainty, authority and temporal continuity can improve agent reliability.

FUNCTIONAL ACCESS

What information becomes globally available and can change behavior?

METACOGNITION

Can the system represent ignorance, confidence and error?

SELF-MODEL

Does it track identity, capability, limits and history?

META-GOVERNANCE

Which authority is governing the decision — and why?

REVERSE CAPTURE

Can the system recover when a partial authority hijacks the whole decision space?

WORLD-EFFECT CLOSURE

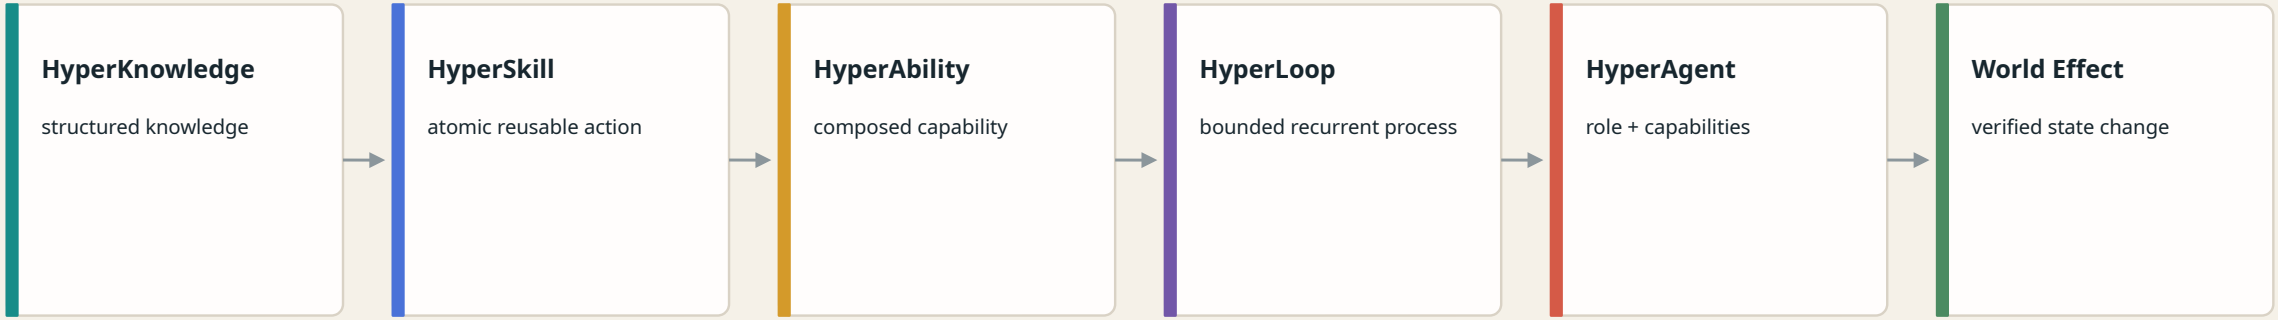
Does judgment survive action, readback and postcondition verification?

Integration: MUT defines the hypotheses and metrics; MUG controls their vessel / authority; MOUSSAIXT provides the test environment.

SOURCE ANCHORS · Governance Space · Reverse Capture · Conscious AGI Master Corpus p.7

Knowledge compiles into governed agency — with a proof path

The framework aims to turn a large corpus into reusable capability without collapsing every idea into weights.



MUT

Defines semantics, relations, failure signatures and tests.

MUG

Decides placement, promotion, authority and rollback.

MOUSSAIXT

Routes, executes, verifies, writes memory and exposes traces.

SOURCE ANCHORS · HyperKnowledge→HyperSkill→HyperAbility→HyperLoop→HyperAgent · Unified AGI

The runtime is the reference implementation target for MUT and MUG

Core layers remain stable; cross-cutting planes add production and evidence capabilities without replacing the cognitive stack.

- L0 · Constitution
- L1 · Mission / Question
- L2 · Δ / MCD / QMCD
- L3–4 · M-HKOS / Evidence
- L5–7 · Skills / Abilities / Loops
- L8–9 · Agents / Tools
- L10 · Courts / Effect
- L11 · Memory / Evolution

CROSS-CUTTING PLANES

- X0 Interoperability
- X1 Durable Execution
- X2 Policy & Trust
- X3 Telemetry & Eval
- X4 Adaptive Memory
- X5 Adaptive Inference
- X6 Semantic Homeostasis
- X7 Capability Certification

X0–X7 are execution planes across L0–L11 — they close interoperability, durability, trust and evidence gaps.

SOURCE ANCHORS · MOUSSAIFXT Cognitive Kernel Part II pp.2–4

The framework has an explicit claim-discipline contract

Canonical theory, external standard, reference implementation and empirical evidence are separate epistemic objects.

SOURCE CONTRACT

CANONICAL [C]
MUT / MUG / internal theory lineage

EXTERNAL [X]
standards · official docs · primary research

SYNTHESIS [S]
reference implementation decisions

E0 → E5

E0 · hypothesis / residual
E1 · object + metric + falsifier
E2 · executable artifact
E3 · local trace + replay
E4 · matched baseline gain
E5 · independent replication

NON-NEGOTIABLE RULES

Architecture ≠ AGI
Code ≠ validated gain
Tool call ≠ world effect
Registry count ≠ scientific evidence
Reported grade ≠ independent grade

Claim Strength ≤ Evidence Strength × Independence × Reproducibility

SOURCE ANCHORS · MUT/MUG Frontier Extension scientific contract · Master Reference E0–E5

What funding should unlock next: the integrated MUT × MUG laboratory

The corpus already supplies candidate mechanisms. The next value is frozen implementation, ablation, capability-placement tests and independent reproduction.

1	**MUT reference bundle**	Freeze 8–12 core mechanisms + baselines + typed metrics
2	**MUG Capability Placement Lab**	Same capability across Context / HyperKB / Skill / Loop / Adapter / Weight
3	**M-HKOS + memory studies**	Retrieval provenance · reconstruction · contradiction retention · writeback
4	**Conscious governance studies**	Meta-governance · reverse capture · uncertainty-to-agency · continuity
5	**MOUSSAIXT runtime closure**	Durable execution · telemetry · world-effect receipts · rollback
6	**E4 → E5 replication**	Matched budgets · frozen baselines · raw traces · independent environments

RESEARCH ASK

Fund the conversion of the corpus into reproducible science.

Compute · research engineering · benchmark design · consciousness studies · replication

Target outcome · one coherent MUT/MUG reference stack with explicit causal tests, reproducible runtime behavior and publishable cooperative evidence.

SOURCE ANCHORS · Master Reference funding map · MUG sponsored study · MUT E4 bundle