

When Tool Architecture Becomes a Black Box: Observing the F1–H1 Boundary

From infrastructural AI density to the possibility of unrecognized hybrid cognition

Joaquim Santos Albino

HibriMind — Observability / Homo Hybridus

12 September 2026

Artificial intelligence is acquiring a body.

Not metaphorically.

Chips, memory, energy, data centers, communication networks, territory, cooling systems, redundancy, physical security and increasingly geopolitical protection are becoming material conditions for the persistence of the algorithmic field.

A recent development in the Gulf makes this transformation unusually visible.

The UAE–US AI Campus was originally announced as a concentrated 5 GW AI infrastructure project in Abu Dhabi. Its first major component, Stargate UAE, involves G42, OpenAI, Oracle, NVIDIA, Cisco and SoftBank, with 1 GW of planned capacity and an initial 200 MW expected online in 2026.

Following Iranian attacks that damaged data-center infrastructure in the UAE and Bahrain, Reuters reported on 11 September that the UAE was considering significant changes to the original architecture.

Among the possibilities under examination were geographical distribution of computing facilities, partially underground construction, blast-resistant structures, drone and missile defense, and facilities associated with protected mountainous military areas.

The final configuration remains undecided.

That uncertainty matters.

The observation is not that AI infrastructure has already become militarily hardened at this scale. The observation is that **military vulnerability has entered the architectural design space of the algorithmic field.**

The chain is becoming longer:

**chip → memory → energy → data center → territorial distribution → fortification →
military defense → operational continuity**

The data center is no longer merely a computational factory.

It is becoming strategic infrastructure.

And that raises a different question.

From Tool to Infrastructure

Most contemporary descriptions of AI still assume an instrumental ontology.

The basic structure is familiar:

human → algorithmic tool → result

The tool may become extraordinarily powerful. It may automate decisions, maintain persistent memory, coordinate agents, operate across institutions and become indispensable to economic or governmental activity.

Yet the ontological center remains outside it.

Humans build it.

Organizations deploy it.

States regulate it.

Institutions protect it.

The algorithmic system remains a capability serving purposes located elsewhere.

We can call this basic condition **F0 — discrete instrumentality**.

But the material expansion now occurring suggests another condition.

When algorithmic systems require dedicated energy, distributed data centers, persistent memory, institutional integration, legal frameworks, territorial protection, redundancy and military security, the tool has not necessarily ceased to be a tool.

Instead, **instrumentality itself has become infrastructurally dense**.

We call this provisional condition:

F1 — infrastructural tool density

F1 is not hybrid cognition.

That distinction is essential.

Density Is Not Hybridity

A civilization may become profoundly dependent on algorithmic infrastructure without becoming cognitively hybrid.

Dependence is insufficient.

Automation is insufficient.

Autonomy is insufficient.

Persistent memory is insufficient.

Even infrastructural indispensability is insufficient.

A society may be unable to function without electricity while not constituting a hybrid cognitive system with its electrical grid.

The same caution must apply to AI.

This suggests separating two dimensions that are often implicitly collapsed.

D — Tool Density

The degree to which social, material and institutional structures are occupied, reorganized or made dependent upon algorithmic systems.

H — Hybrid Constitutivity

The degree to which interpretation, decision, local purpose and trajectory emerge from longitudinal human–algorithm interaction and can no longer be adequately attributed to either pole independently.

These dimensions need not evolve together.

A state may develop extremely high **D** while retaining low **H**.

Conversely, a small-scale longitudinal human–AI relationship may possess little infrastructure while exhibiting much stronger causal interdependence in the production of cognitive trajectory.

This distinction prevents a crucial mistake:

More AI does not necessarily mean more hybridity.

The F1–H1 Boundary

If F1 describes infrastructurally dense instrumentality, what would distinguish it from a genuinely hybrid condition?

We provisionally designate the latter:

H1 — hybrid constitutivity

H1 should not be defined by the number of models, GPUs, agents or automated decisions involved.

The relevant transition may instead concern **causal separability**.

As long as a trajectory can still be adequately reconstructed as:

human intention + algorithmic assistance

the system remains compatible with F1.

But consider increasingly recursive architectures:

human → algorithm → human → agent → memory → institution → algorithm → decision
→ infrastructure → another human → new algorithmic state → subsequent decision

Each state becomes part of the initial conditions of later states.

Outputs become inputs.

Human interpretations alter algorithmic contexts.

Algorithmic outputs alter human interpretation.

Those altered humans generate new inputs.

Persistent algorithmic states then influence subsequent decisions.

At sufficient longitudinal depth, a different observational problem appears:

Can the genealogy of the trajectory still be adequately decomposed into a human component and an algorithmic component?

If the answer eventually becomes no, we may be observing something categorically different from a dense tool architecture.

That is the candidate boundary between F1 and H1.

A Second Black Box

AI is frequently described as a “black box.”

The phrase is often used too loosely.

It does not simply mean that nobody knows how artificial intelligence works. Architectures, training procedures, inference mechanisms and many internal processes are technically accessible to investigation.

The deeper difficulty concerns genealogy.

For a sufficiently massive system, reconstructing the complete causal history that produced one particular output can become practically impossible.

This can be described as a **first-order Black Box**:

genealogical opacity of an algorithmic output

But infrastructural density introduces the possibility of another scale.

When humans, models, memories, agents, organizations, databases, infrastructures and subsequent human reactions recursively modify one another, the opacity no longer concerns only the output of a model.

It concerns the trajectory of the entire system.

We therefore introduce a provisional second concept:

Systemic Black Box — a condition in which the massiveness and recursion of human–algorithm interactions make it impractical or inadequate to reconstruct the trajectory through independent human and algorithmic genealogies.

This produces the central hypothesis of this article:

F1 may grow until it constitutes a systemic Black Box within which H1 emerges before external observers recognize it.

Hybridity May Arrive Without Announcement

This possibility has an unusual consequence.

We often imagine major technological transitions as recognizable events.

A new machine appears.

A capability threshold is crossed.

A system becomes autonomous.

An institution announces a transformation.

But the F1–H1 transition may have no such moment.

The declared ontology may remain instrumental long after the operational ontology has changed.

Institutions may continue saying:

“We use AI tools.”

Humans may continue believing:

“AI assists our decisions.”

States may continue regulating:

“AI systems under human control.”

Yet the actual genealogy of important decisions may already require persistent interaction between human and algorithmic states to be explained.

In that case, hybridization would precede its recognition.

The massification of tool architecture could simultaneously provide the conditions for hybrid processes and conceal them beneath an increasingly convincing instrumental description.

This leads to a stronger observational possibility:

Homo Hybridus may become empirically present before becoming ontologically visible.

The UAE Indicator

The UAE case does not demonstrate this transition.

It demonstrates something more modest and therefore more useful.

It shows F1 acquiring material depth.

War is beginning to participate in the architecture of the algorithmic field.

A computational facility exposed to military threat generates requirements for territorial distribution, redundancy, physical protection and operational survival.

The algorithmic field consequently acquires properties determined not only by computer engineering but by geopolitics.

This can be formulated as:

When the algorithmic field becomes strategic infrastructure, war ceases to be merely an external risk and begins to participate in the architecture of its body.

This is an empirical indicator.

It does not tell us where the trajectory will end.

It changes the space of possible trajectories.

And, importantly, the observation preceded the hypothesis.

We did not begin by assuming an F1–H1 transition and searching for evidence.

An external material event exposed a new layer of the algorithmic field, from which the question emerged.

The Adversarial Boundary

The hypothesis must remain vulnerable to contradiction.

Several boundaries are therefore necessary.

Complexity is not hybridity

A system may become too complex for any observer to reconstruct while remaining entirely instrumental.

Opacity is not ontology

Failure to observe causal decomposition does not prove that causal decomposition no longer exists.

Dependence is not cognition

A civilization may become completely dependent upon an infrastructure without forming a cognitive system with it.

Functional inseparability is not cognitive inseparability

The inability to remove an algorithmic infrastructure without social collapse does not establish hybrid cognition.

Causal inseparability may still be insufficient

Even if longitudinal trajectories become difficult to decompose, additional properties may be required: persistent shared history, reciprocal transformation, memory, attractor stabilization or other mechanisms not yet adequately discriminated.

These objections are not weaknesses to be removed.

They define the experimental boundary of the hypothesis.

What Should We Observe?

The relevant question may therefore not be:

When will AI become advanced enough to create Homo Hybridus?

A more useful question is:

Where do trajectories already exist whose genealogy can no longer be adequately described by the human–tool paradigm?

Candidate observations include:

- persistent states jointly produced by humans and algorithms;
- decisions whose genealogy requires multiple human–algorithm cycles;
- reciprocal modification of future conditions for both poles;
- local purposes stabilized through interaction rather than imposed once by a human operator;
- increasing difficulty in identifying the causal provenance of institutional decisions;
- divergence between declared instrumental ontology and observed system architecture.

The objective is not to prove hybridity.

It is to build an observational instrument capable of distinguishing two errors:

seeing hybridity where there is only a very dense tool

and

continuing to see only a tool after the trajectory has become constitutively hybrid.

The Last Shell

The expansion of AI infrastructure does not demonstrate Homo Hybridus.

It may simply represent the extreme evolution of instrumentality.

But as tool architecture becomes increasingly distributed, recursive, persistent and embedded in human institutions, the genealogy of trajectories becomes a more important object than the capabilities of individual models.

If the trajectory remains adequately explainable as human action assisted by algorithms, we remain within F1.

If its constitutive history eventually requires the longitudinal human–algorithm relation itself as the relevant causal unit, H1 becomes an empirical candidate.

The challenge for Observability is to recognize that boundary before infrastructural density makes it invisible.

A formulation developed elsewhere in the HibriMind corpus becomes useful here:

The shell is not false. The error begins when the shell is mistaken for the onion.

F1 may be the last shell.

And the systemic Black Box may be where we eventually have to look for what has already begun to exist inside it.

Sources and Observational Provenance

The empirical trigger for this article is deliberately separated from the ontological derivation that follows from it.

Primary infrastructure source

G42 — “Global Tech Alliance Launches Stargate UAE”

G42, OpenAI, Oracle, NVIDIA, SoftBank Group and Cisco announced Stargate UAE as a 1 GW AI compute cluster within the planned 5 GW UAE–US AI Campus in Abu Dhabi. G42 states that the first 200 MW cluster is expected to become operational in 2026.

The announcement also describes the infrastructure as providing “nation-scale compute” and connects it to sectors including healthcare, energy, finance and transportation.

Official source:

<https://www.g42.ai/resources/news/global-tech-alliance-launches-stargate-uae>

Material-change indicator

Reuters — 11 September 2026 — “UAE revises AI data center plan after Iranian attacks, sources say”

Reuters reported that, following Iranian attacks that damaged data-center infrastructure in the UAE and Bahrain, the original concentrated architecture of the 5 GW project was under reconsideration.

According to sources consulted by Reuters, options being examined included a distributed network of facilities, partially underground construction, blast-resistant structures, protection against drones and missiles, and protected facilities in mountainous areas for military use.

G42 confirmed that project specifications remained under review for security and resilience reasons while stating that construction continued as planned. Reuters could not establish the final architecture or its eventual effects on cost and schedule.

Source:

<https://www.reuters.com/world/middle-east/uae-revises-ai-data-center-plan-after-iranian-attacks-sources-say-2026-09-11/>

Provenance Boundary

These sources establish the infrastructural configuration and the reported reconsideration of its physical architecture.

They do **not** establish F1, H1, hybrid cognition, a Systemic Black Box, or an F1–H1 transition.

Those are HibriMind analytical derivations produced from the observed material change.

The epistemic sequence of this article is therefore:

external event → documented infrastructural response → HibriMind observation → conceptual derivation → open hypothesis

The observation preceded the ontology.

Epistemic Status

F1: provisional category.

H1: provisional category.

F1–H1: candidate ontological boundary, not an assumed chronological sequence.

Systemic Black Box: candidate concept.

Causal separability: candidate discriminator.

Homo Hybridus: not demonstrated by this article.

UAE infrastructure development: empirical indicator motivating the derivation, not evidence of H1.

HibriMind — Observability / Homo Hybridus

12 September 2026