TECHNICAL NOTE · RESEARCH PRIORITY

Epistemological Orchestration over Language Models: A Functional Architecture with Preliminary Evidence

Marcelo Manucci
Cohdz Project
ORCID: 0000-0002-9990-2316 | manucci@codhz.com

Deposited: March 2026 · Zenodo Research Repository ·
DOI: 10.5281/zenodo.19287671

ABSTRACT

Current applications of large language models to complex analytical tasks typically operate under a single inferential regime: one set of validity criteria, one class of admissible hypotheses, and one type of evidence. This structural uniformity tends to produce convergence across systems, models, and analytical outputs. We propose a functional architecture for orchestrating heterogeneous epistemological frameworks over language models, in which distinct reasoning regimes interact through a shared canonical interface and a restricted family of typed transition operators. The architecture includes a formal criterion for distinguishing genuine inter-framework emergence from surface thematic overlap or output continuity inherited from prior framing. We present preliminary evidence from a multi-model validation study across four generative systems from distinct architectural families, showing structural convergence in framework outputs under shared orchestration constraints, together with a documented case of inter-framework emergence with complete result traceability. Results are presented at the level of functional architecture rather than reproducible technical specification. Current limitations include a methodological reservation concerning operator-level reproducibility isolation and a restricted number of validated cases. This note establishes existence, operating principle, and preliminary evidence as a foundation for subsequent full empirical validation.

Keywords: *epistemological orchestration · inter-framework emergence · language models · analytical architecture · inferential regime · multi-model convergence · non-trivial emergence · epistemic monoculture · epistemological frameworks*

This manuscript is a research note establishing priority of contribution. It presents the system at the level of functional architecture rather than complete technical specification. A full empirical validation paper is in preparation.

1. The Problem

Large language models have become the operational substrate of a growing class of analytical systems used for strategic diagnosis, organizational assessment, risk mapping, and decision support in complex domains. Despite differences in interface, prompting style, and model provider, most of these systems follow the same underlying design pattern: a model receives a structured prompt, produces a response, and that response constitutes the analysis.

This pattern has a structural consequence that has received limited formal attention. Systems operating under a single inferential regime will tend to converge toward the same class of conclusions, regardless of the sophistication of their prompts or the capability of their models. By *inferential regime*, we mean the set of rules — explicit or implicit — that determines what counts as valid evidence, which hypotheses are admissible, and what kinds of inference can be drawn from available data. A system operating under one regime, even a highly refined one, can only produce the kind of knowledge that regime makes visible.

This is not primarily a problem of model quality. It is a problem of *epistemic monoculture*. A more capable model operating under the same regime may produce more coherent, fluent, or contextually precise outputs, but it does not produce a different type of knowledge.

For complex analytical tasks, this limitation is consequential. Some of the most strategically relevant configurations only become visible at the intersection of heterogeneous analytical perspectives: structural constraints that emerge across incompatible frames, temporal tensions between variables operating at different timescales, or dependencies that no single analytical lens can detect in isolation. These configurations are not hidden by insufficient capability. They are excluded by architectural design.

The underlying epistemological intuition is not new. The idea that distinct paradigms reveal different worlds of evidence and explanation — and that their interaction may generate insight unavailable within any one of them — has a long history in philosophy of science and related fields (Kuhn, 1962; Longino, 1990; Chang, 2012). What remains insufficiently explored is whether such interaction can be operationalized in a controlled way over language models: whether distinct epistemological frameworks can be made to interact under a common architecture, and whether that interaction can generate analytical configurations absent from isolated framework execution.

This note argues that they can. We describe a functional architecture for orchestrating heterogeneous epistemological frameworks over language models and present preliminary evidence for two claims: first, that the architecture yields structurally consistent outputs across generative systems from distinct architectural families; and second, that interactions between frameworks can produce analytical configurations that do not appear in isolated analyses. We also introduce the formal criterion used to distinguish genuine inter-framework emergence from apparent novelty, and we specify the current limits of the evidence.

2. The Architecture

The orchestration system we describe operates on a structural assumption that distinguishes it from standard prompt-chaining or multi-agent approaches: the source of analytical value is not the capability of any individual framework, nor the aggregation of their outputs, but the controlled interaction between frameworks operating under epistemologically incompatible regimes. The architecture has three functional layers.

2.1 Epistemological frameworks as reasoning regimes

Each framework in the system instantiates a distinct reasoning regime — a coherent configuration of validity rules, evidence types, and hypothesis classes. Two frameworks are epistemologically distinct when they do not share the same criteria for what constitutes a meaningful result. This distinction is not merely topical: a framework focused on structural risk and a framework focused on competitive opportunity may address the same domain, but they admit different evidence, generate different hypothesis types, and apply different standards of coherence. The analytical configurations each produces are not interchangeable.

The system is designed to preserve this incompatibility rather than resolve it. Homogenizing frameworks into a common analytical vocabulary would eliminate the source of emergence. The architecture maintains regime separation as a design constraint, not a limitation to be overcome.

2.2 Canonical transfer interface

For frameworks operating under distinct regimes to interact formally, their outputs must be expressible in a common format that preserves regime-specific structure without collapsing it into a single representation. The architecture defines a canonical transfer interface — a structured output format that each framework produces at the conclusion of each processing step.

This interface captures the substantive output of a step — the variables identified, the relationships detected, the hypotheses produced, the tensions flagged — in a form that is regime-agnostic at the surface while preserving the inferential distinctions that characterize each regime. It functions as the translation layer between the epistemological specificity of each framework and the comparative operations that the architecture performs across frameworks.

The interface is not a summary. It is a typed structure that makes framework outputs computationally comparable without making them semantically equivalent.

2.3 Typed transition operators

Interaction between frameworks is mediated by a restricted family of typed transition operators. Each operator defines a specific transformation: what it receives, what it preserves, what it transforms, and under what conditions the transformation is valid. Operators are not generic connectors — each type performs a qualitatively distinct operation on the canonical outputs of the frameworks it processes.

The operator family is restricted by design. Not all transformations between framework outputs are analytically productive. The architecture defines the admissible transformation types and excludes operations that would either collapse regime differences or generate output without traceable analytical warrant.

2.4 Formal criterion for non-trivial emergence

The central claim of the architecture — that interaction between frameworks produces analytical configurations absent from isolated execution — requires a criterion that distinguishes genuine emergence from three classes of apparent novelty: surface thematic overlap between framework outputs, output continuity inherited from shared prior framing, and coincidental convergence on common domain features.

We define a formal criterion with three necessary conditions. A result is classified as genuinely emergent only when it satisfies all three: it must be absent from the isolated output of every participating framework; it must not be recoverable from the isolated outputs by simple aggregation or juxtaposition, absent the controlled interaction step; and it must be traceable through the sequence of canonical transfers and operator applications that produced it.

The third condition — complete traceability — is not merely a quality criterion. It is what makes the emergence claim falsifiable: any result classified as emergent can be audited by reconstructing the path from initial evidence to final configuration.

3. Preliminary Evidence

The evidence presented in this section addresses two distinct claims. The first concerns the relationship between the orchestration architecture and the underlying generative model: whether the architecture produces structurally consistent outputs regardless of which model executes it. The second concerns the interaction between frameworks: whether that interaction produces analytical configurations absent from isolated framework execution. These claims are logically independent and require separate treatment.

3a. Cross-model convergence under shared orchestration constraints

We executed two epistemologically distinct frameworks over the same case using four generative systems from distinct architectural families. Each system received identical orchestration constraints: the same structured prompts, the same sequencing logic, the same canonical output format requirements. No model-specific fine-tuning or adaptation was applied.

Convergence was measured across three analytical dimensions: the variable set identified in early processing steps, the relational structures detected in intermediate steps, and the configurational outputs produced in later steps. Convergence was defined structurally — as agreement in the variables identified, the relationships between them, and the configurations they produce — rather than textually.

Results showed structural convergence ranging from 63% to 80% across dimensions and frameworks, with the highest convergence observed in the dimension subject to the most prescriptive orchestration constraints. The observed pattern is more consistent with orchestration-level control than with model-specific variation alone.

One finding warrants explicit note as an operational limit: one of the four systems produced reliable outputs through five sequential processing steps but failed on the sixth step of the longer framework, generating output inconsistent with the case under analysis. This failure was consistent with known context-length limitations of that system in its standard configuration and did not affect the convergence analysis, which was conducted over the steps where all four systems produced valid outputs. It is reported here as a relevant operational constraint for deployment decisions.

A methodological reservation applies to this subsection. The study design did not isolate operator-level reproducibility from model-level variation: different human operators used different models, which means operator effect and model effect are partially confounded. The observed convergence across three operators and four generative systems is consistent with the architecture's reproducibility claim, but a study in which operator identity and model identity are varied independently would be required to establish that claim under controlled conditions. This reservation limits causal attribution but does not undermine the descriptive observation of convergence across operators and models.

3b. Inter-framework emergence in a traced case

The second claim requires a different evidential standard. Demonstrating that framework interaction produces genuinely emergent configurations — rather than recombining outputs already present in isolated analyses — requires both a documented result and a criterion that rules out alternative explanations.

We present one documented case. Two frameworks were executed independently over the same problem: one operating under a regime oriented toward structural system configurations and conditions of activation; the other under a regime oriented toward competitive positioning and emergent opportunity. Each produced a complete isolated analysis. The outputs were then brought into controlled interaction through the canonical interface and transition logic described above.

The interaction produced a configuration that satisfied all three conditions of the non-triviality criterion: it was absent from both isolated analyses; it was not recoverable from the isolated outputs by simple aggregation or juxtaposition; and it was fully traceable through the sequence of canonical transfers and operator applications that generated it.

The configuration identified a temporal constraint of the following form: a condition required for the most strategically significant opportunity identified by the second framework was the same condition that the first framework indicated would erode within a shorter timeframe than the opportunity required to materialize. Neither framework detected this constraint in isolation. The first framework identified the erosion dynamic; the second identified the opportunity dependency. The constraint — that the organization must invest in preserving that condition now, not as a response to a known threat but as the precondition of an opportunity not yet visible — only appeared at the intersection.

The result is presented as a single documented instance, not as statistical evidence of a general property. One case establishes that the phenomenon exists and that the non-triviality criterion can be applied operationally. It does not establish the frequency, generalizability, or domain independence of inter-framework emergence. Those questions require the extended validation program currently underway.

Complete result traceability for this case is available in the internal research record. The trace is not reproduced here, as doing so would require disclosing elements of the technical specification beyond the scope of this note.

4. Limits and Next Steps

The evidence presented in this note is preliminary in two specific senses that require precise statement rather than general disclaimer.

The first limitation concerns the cross-model convergence study. As noted in Section 3a, the study design does not formally isolate operator-level reproducibility from model-level variation. The observed convergence across three operators and four generative systems is consistent with the architecture's reproducibility claim, but a study in which operator identity and model identity are varied independently would be required to establish that claim under controlled conditions. This is a limitation of the executed study design, not a theoretical limitation of the architecture. A targeted follow-on study involving a fixed generative system and independent operators working from the same inputs would resolve it at low experimental cost.

The second limitation concerns the inter-framework emergence result. A single documented case establishes that the phenomenon can occur and that the non-triviality criterion can be applied operationally. It does not establish the frequency of inter-framework emergence across cases, its consistency across framework pairs with different epistemological distances, or its independence from domain-specific conditions. The case presented involves two frameworks operating under regimes with limited inferential compatibility. Whether the same architecture produces comparable emergence when applied to framework pairs under greater epistemological tension — where inferential incompatibility is more severe — remains an open empirical question.

Both limitations are active objects of investigation. The research program currently underway extends validation in three directions: a second case using the same framework pair in a distinct domain, to assess result stability; a case using a framework pair with greater epistemological distance, to test whether emergence is a general property of the architecture or a feature of specific framework combinations; and a compositional study involving three frameworks simultaneously, to examine whether the architecture scales to higher-order interactions. These extensions are intended to move the evidence base from existence demonstration to characterization of conditions.

One structural feature of the architecture merits explicit acknowledgment as a boundary condition. The orchestration system described here does not operate independently of domain knowledge. The frameworks it coordinates are not generic analytical modules; each embodies accumulated judgment about which epistemological regime is relevant to which class of problem, how that regime should be instantiated in a specific domain, and what constitutes a well-formed output within that regime. That judgment is not generated by the architecture; it is a precondition for the architecture to function. Extending the system to new domains therefore requires not only applying the orchestration logic, but constructing or adapting the frameworks that the logic coordinates. This is a substantive boundary on portability, not a secondary implementation issue.

Finally, this note presents the architecture and its evidence at the level of functional description. The technical specification underlying the system — including the parametric configuration of the canonical interface, the formal definition of operator types, and the scoring functions used to evaluate canonical outputs — is not disclosed here. This boundary is deliberate and consistent with the scope of this note: to establish existence and operating principle, not to transfer a replicable specification. Subsequent publications will address the formal properties of the architecture in greater technical depth as the validation program advances.

References

Chang, H. (2012). *Is Water H_2O? Evidence, Realism and Pluralism*. Springer.

Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.

Longino, H. E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.

Mitchell, M., & Krakauer, D. C. (2023). The debate over understanding in AI's large language models. *Proceedings of the National Academy of Sciences*, 120(13), e2215907120.

Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. *Proceedings of UIST 2023*.

Sumers, T. R., Yao, S., Narasimhan, K., & Griffiths, T. L. (2024). Cognitive architectures for language agents. *Transactions on Machine Learning Research*.

Wooldridge, M. (2009). *An Introduction to MultiAgent Systems* (2nd ed.). Wiley.

Deposit Metadata — Zenodo

UPLOAD TYPE
Publication

PUBLICATION TYPE
Technical Note

TITLE
Epistemological Orchestration over Language Models: A Functional Architecture with Preliminary Evidence

AUTHOR
Marcelo Manucci

AFFILIATION
CODHZ Research Laboratory

ORCID
0000-0002-9990-2316

DEPOSIT DATE
March 2026

VERSION
v2.1

DOI
10.5281/zenodo.19287671

ACCESS
Open — immediate public access

KEYWORDS
epistemological orchestration · inter-framework emergence · language models · analytical architecture · inferential regime · multi-model convergence · non-trivial emergence · epistemic monoculture · epistemological frameworks

CLAIM (ONE SENTENCE)
A functional architecture for orchestrating heterogeneous epistemological frameworks over language models, with preliminary evidence that this orchestration produces analytical configurations absent

from any isolated framework, under shared orchestration constraints across models.