

Agemic AI vs. Agentic AI: Toward an Agemics of Artificial Agency

Dmitry Vostokov (DumpAnalysis.org)
with GPT-5.4 (AI-assisted drafting)

Agemic AI vs. Agentic AI

Toward an Agemics of Artificial Agency

AGENTIC AI The level of realized agency. It concerns systems that plan, act, use tools, monitor, recover, delegate, and complete tasks.	AGEMIC AI The level of agency structure. It concerns the minimal functionally contrastive units of agency—agemes—from which agentic behavior is composed.
--	--

FROM SURFACE ACTIONS TO STRUCTURE THAT MATTERS
Agemic analysis transforms agentic traces into a structurally interpretable language of agency.

MINIMAL & CONTRASTIVE Agemes are the minimal functionally contrastive units of agency.	REALIZATIONAL PLURALITY Each ageme can be realized in multiple ways across systems.	COMPOSITIONAL Agemes combine to form stable patterns of artificial agency.	DIAGNOSTIC POWER Reveals absent, weak, conflicting, or miscomposed units of agency.	COMPARE & GENERALIZE Enables agemic equivalence across different implementations.
--	---	--	---	---

Agentic AI builds systems that act. *Agemic AI helps us understand the structure of why they do.*

Table of Contents

Table of Contents	1
Abstract	2
Introduction	3
Why a New Distinction is Needed	3
The Phonetics and Phonemics Analogy	4

Agentic AI as Realized Artificial Agency	5
Agemic AI as the Study of Agency Structure	5
Agemes and Contrastive Function	6
Criteria for Individuating Agemes	7
Agentic Tokens and Agemic Types	7
Agemics as a Field	8
Agemic Transcription	8
Agentic and Agemic Equivalence	10
Relation to Concrete and General Levels in Pattern-Oriented Diagnostics	11
A Pattern-Oriented Diagnostics Interpretation	12
Agemic Diagnostics of Failure	12
Agemic Repertoire and Minimal Agency	13
Phonetic Variation, Allophony, and Surface Realization	13
Surface Agency and Deep Agency Structure	14
Illustrative Comparison	14
Originality and Conceptual Novelty	14
Implications for Future Work	15
Conclusion	16
Vocabulary	17

Abstract

Agentic AI is usually discussed at the level of whole systems: planning, tool use, memory, delegation, monitoring, and task completion. Such discussion remains too coarse unless accompanied by a second analytical level, which the article introduces as *Agemic AI*. By analogy with the distinction between phonetics and phonemics, *Agemic AI* concerns the *minimal functionally contrastive units of agency* from which realized agentic behavior is composed. These units, here called *agemes*, are not defined merely as important steps or implementation modules, but as structurally contrastive units satisfying criteria of contrast, realizational plurality, minimality, compositionality, and diagnostic yield. On this basis, the article proposes a method for reducing concrete agentic traces to agemic transcriptions and for distinguishing agentic similarity from agemic equivalence. The paper's contribution is therefore not merely terminological. It is a criterion-guided framework for the structural analysis and diagnosis of artificial agency.

Introduction

The expression *Agentic AI* has become increasingly prominent in discussions of contemporary artificial intelligence. It names systems that do more than produce isolated responses to isolated prompts. Such systems appear to act. They maintain goals across steps, select means, retrieve information, use tools, adapt to intermediate feedback, coordinate subprocesses or subagents, and persist in a task until some stopping condition or completion criterion is reached. In this sense, *Agentic AI* designates AI at the level of the agent.

This level is important and practical. It is the level at which systems are engineered, evaluated, deployed, and discussed in applications. Yet it may also be conceptually too coarse. It tells us that a system behaves agentially, but not what the minimal meaningful constituents of its agency are. It identifies organized purposive behavior, but it does not isolate the elementary structural units from which such behavior is composed. It characterizes agency as a whole, but not its structure.

For this reason, it is useful to introduce a parallel and more abstract notion: *Agemic AI*.

Agemic AI is a conceptual invention modeled on the distinction between phonetics¹ and phonemics², or, more broadly, between phonetic realization and phonological structure. Just as phonetics studies the concrete realization of speech sounds while phonemics studies the minimal functionally contrastive sound units that matter structurally in a language, so *Agentic AI* may be understood as the study of concrete realized agency in artificial systems, whereas *Agemic AI* may be understood as the study of the minimal functionally contrastive units of agency from which such realized behavior is composed.

The central claim of this article is therefore not merely that agency can be decomposed, but that it can be decomposed in a disciplined way into *minimal functionally contrastive units* whose identification permits reduction, comparison, and diagnosis of agentic behavior. *Agemic AI* is proposed as the level at which such units are analyzed. The aim is to move from surface descriptions of agentic workflows to a more precise structural account of what units of agency are being instantiated, how they combine, and where they fail.

Why a New Distinction is Needed

Contemporary discourse about *Agentic AI* often remains at the level of visible system properties. A model or system is called agentic when it plans, calls tools, uses memory, collaborates with other components, or carries out multi-step tasks. This usage is reasonable, but it frequently conflates concrete realization with structural organization.

¹ <https://en.wikipedia.org/wiki/Phonetics>

² <https://en.wikipedia.org/wiki/Phonology>

Two systems may both appear agentic while differing profoundly in how they organize agency. One may genuinely monitor itself and revise plans in light of failure, while another may only simulate sequentiality without a stable repair mechanism. Conversely, two implementations may differ in code, architecture, prompting style, or orchestration platform while realizing similar agency structures.

A more abstract descriptive level is therefore needed. Such a level should abstract away from implementation details without losing the distinctions that matter structurally. It should preserve functionally significant units while ignoring incidental variation. The analogy with phonetics and phonemics provides exactly such a model.

The Phonetics and Phonemics Analogy

The analogy with phonetics and phonemics is central here. The analogy is structural rather than literal: it proposes a comparable level distinction, not a one-to-one identity between linguistic and agential phenomena.

Phonetics concerns the concrete realization of speech sounds. It studies articulation, acoustics, and auditory form. The phonetic level is materially rich, continuous, and finely differentiated. Many phonetic differences can be observed and measured even when they do not matter structurally in a given language.

Phonemics, by contrast, concerns the sound distinctions that are structurally significant. A phoneme³ is not merely any sound. It is a functionally contrastive unit within a system. Different phonetic realizations may count as instances of the same phoneme if they do not alter linguistic function, while a relatively small phonetic difference may be phonemically decisive if it changes meaning or grammatical organization.

This relation can be transferred fruitfully to artificial agency.

At the agentic level, one observes the concrete realization of agency in a system. One sees prompts, planning loops, retrieval chains, tool calls, memory reads and writes, delegations, critiques, retries, corrections, and final answers. This corresponds to the phonetic level of actual realization.

At the agemic level, one abstracts from concrete realization and asks which distinctions are functionally and structurally relevant for agency. One no longer focuses on every observed step as such, but on the minimal units whose presence, absence, substitution, or recombination changes the organization of agency itself.

Thus, the analogy may be stated as follows. Phonetics is to speech realization as Agentic AI is to agency realization. Phonemics is to contrastive sound structure as Agemic AI is to contrastive agency structure.

³ <https://en.wikipedia.org/wiki/Phoneme>

This is not merely a decorative metaphor. It is a methodological proposal. The task of Agemic AI is to identify the agency-level analog of phonemic structure.

The term Agemic AI is grounded primarily in the analogy with phonemics, since its main concern is with minimal functionally contrastive units of agency. At the same time, the suffix also resonates with the broader emic tradition in anthropology and semiotics, which aims to identify categories meaningful within the organization under study rather than merely convenient for external description. In this sense, agemic analysis is not simply an experience-near surface description of agentic traces, nor merely an experience-far theoretical redescription. It seeks a second-order structural account of agency that remains faithful to intrinsic contrasts within agentic organization itself.

Agentic AI as Realized Artificial Agency

Agentic AI refers to AI systems that behave as agents. Such systems typically do some combination of the following. They receive or formulate goals. They decompose tasks into subtasks. They search for information. They select and invoke tools. They monitor intermediate results. They revise plans in response to feedback. They recover from errors. They coordinate multiple components or subagents. They maintain context or memory across steps. They produce outputs in light of evolving internal and external states.

What makes such systems agentic is not mere intelligence in the narrow sense, nor mere fluency of output, but organized purposive behavior across time. Agentic systems are not simply generators of isolated continuations. They appear to sustain an action structure.

This level is practical and operational. It is the level at which one designs systems, inspects workflows, studies observability artifacts, and evaluates performance on tasks. Yet it remains the level of realized wholes. It tells us that something acts like an agent, but it does not yet isolate the elementary units that make such agency possible.

Agemic AI as the Study of Agency Structure

Agemic AI concerns the minimal functionally contrastive units of agency.

An *ageme* may be defined as the smallest functionally contrastive unit of agency in an AI system. The term is constructed by analogy with phoneme. Just as a phoneme is not simply any sound whatsoever, an *ageme* is not any event, module, or line of code. It must be a unit whose presence, absence, or substitution significantly alters the structure of agency.

An *ageme* is therefore not merely computational or operational. It is structural and functional. It belongs not to raw execution detail but to the organization of agency itself.

Candidate agemes may include goal commitment, plan formation, action selection, retrieval, tool invocation, environmental perception, memory update, self-monitoring, error recovery, delegation, closure, and termination. These are not necessarily implementation modules. They are abstract units in a repertoire of agency. A concrete API call is not itself an ageme simply because it occurred. Rather, it may instantiate an agemic unit such as external information acquisition or tool-mediated action.

The status of these candidates is not automatic. A proposed ageme must be justified by contrastive function rather than by intuitive importance alone. Goal commitment, retrieval, monitoring, repair, delegation, closure, and termination are therefore offered not as dogmatic primitives, but as candidates whose agemic status depends on whether they satisfy criteria of contrast, realizational plurality, minimality, compositionality, and diagnostic yield. Agemic AI thus concerns not the arbitrary naming of agent functions, but the disciplined identification of structurally significant units of agency.

Agemes and Contrastive Function

The strongest way to understand agemes is not through size but through contrastive function.

A unit is agemic not because it is small in a physical or code-level sense, but because it is structurally decisive. If removing or replacing it changes what kind of agency the system can exhibit, then it is a candidate ageme.

Consider two systems that both generate answers, retrieve documents, and call tools. On the surface, they may appear similar. Yet one may possess a genuine repair mechanism that re-evaluates failed results and revises strategy, while the other merely continues linearly without effective correction. In agentic terms, both may look active. In agemic terms, they differ. One system contains a repair ageme in a structurally meaningful form; the other may not.

Conversely, two systems may differ greatly in codebase, architecture, prompt design, or model family and yet remain agemically equivalent if the same functional units of agency are preserved. One may implement delegation through explicit orchestration, another through prompt-mediated role distribution, but both may instantiate structurally similar delegation.

This is the analytic gain of Agemic AI. It permits comparison across implementations by abstracting away from incidental realization while retaining organizational significance.

Criteria for Individuating Agemes

An ageme should not be treated merely as an important step, module, or capability of an agentic system. To qualify as an ageme, a candidate unit should satisfy several criteria.

First, it should satisfy a *contrast criterion*. Its presence, absence, or substitution should alter the organization of agency rather than merely alter implementation detail. A difference that leaves the structural role unchanged is not yet agemic.

Second, it should satisfy a *realizational plurality criterion*. The same ageme should be capable of more than one concrete realization. Otherwise, the proposed unit risks collapsing into a specific mechanism, module, or tool choice rather than an abstract unit of agency.

Third, it should satisfy a *minimality criterion*. The candidate should not be decomposable into smaller units that preserve the same contrastive role. If such decomposition is possible without loss of structural significance, the candidate is not yet minimal.

Fourth, it should satisfy a *compositional criterion*. An ageme should be capable of entering into multiple larger agency patterns rather than appearing as a single fixed stage in a single workflow. Agemes are intended as reusable structural units of agency.

Fifth, it should satisfy a *diagnostic yield criterion*. Describing an agentic episode in terms of the proposed ageme should improve the explanation of success, failure, instability, or miscomposition relative to ordinary surface workflow descriptions.

Sixth, it may also satisfy an *oppositional criterion*. A candidate ageme is stronger if it participates in a meaningful system of structural opposition, completion, or asymmetry whose disruption has diagnostic significance.

Taken together, these criteria distinguish agemes from mere implementation components, surface actions, or convenient labels. They make Agemic AI not merely a vocabulary of agency, but a disciplined method for identifying structurally significant units of agency.

Agentic Tokens and Agemic Types

The distinction between agentic and agemic levels may also be described as a distinction between token and type.

At the agentic level, one finds concrete occurrences in time. A system may issue a search query, call an API, open a browser, ask another model for help, retry a failed computation, or summarize retrieved text. These are token events of realized agency.

At the agemic level, one identifies the types of agency those token events instantiate. A search query may instantiate information acquisition. A calculator call may instantiate tool-mediated execution. A retry may instantiate repair. A request to another model may instantiate delegated reasoning. A final answer may instantiate closure.

Thus, an execution trace may be richly agentic and yet reducible to a more abstract agemic description. Just as a stream of utterance may be analyzed phonemically, a stream of agentic events may be analyzed agemically.

Agemics as a Field

If ageme is the unit, then *Agemics* may designate the field that studies agemes and their combinations.

Agemics asks questions such as these. What is the agemic repertoire of a given system? Which agemes are necessary for minimal agency? Which agemic oppositions distinguish one family of systems from another? How do agemes combine into stable patterns of action? What kinds of failures correspond to agemic absence, weakness, conflict, oscillation, or over-expansion?

In this sense, Agemics is not merely a terminology. It is the study of structural units of agency, their compositions, and their diagnostic significance.

Agemic Transcription

Just as phonemic transcription reduces phonetic variation and preserves only the distinctions that matter structurally, agemic transcription aims to reduce the surface detail of an agent's behavior and preserve only the units that matter for agency structure. The reduction is not merely a renaming of steps. It is a structural interpretation of an agentic episode in terms of agemes.

A synthetic but operationally realistic example may clarify the distinction. Consider two retrieval-augmented support agents asked the same question: whether the organization's EU data-retention policy permits keeping failed-login IP addresses for more than 90 days.

Both agents have access to a policy index, a keyword search tool, and a summarization tool. At the surface level, both appear agentic. Both retrieve documents, produce intermediate reasoning, and return answers. Yet their internal organization of agency differs.

A simplified trace for Agent A may be represented as follows:

```
[1] receive_task("EU data-retention policy for failed-login IP addresses")
[2] decompose_goal("find policy", "locate retention clause", "check exception")
[3] search_policy_index("failed-login IP retention")
```

```

[4] retrieve_document("Security Logging Standard v3.2")
[5] inspect_clause("retain failed-login metadata for 90 days")
[6] search_policy_index("retention exception security investigation")
[7] retrieve_document("Incident Response Exception Policy")
[8] inspect_clause("security investigations may extend retention with approval")
[9] compare_clauses()
[10] detect_missing_condition("approval requirement not yet verified")
[11] retrieve_document("Approval Workflow for Extended Retention")
[12] inspect_clause("DPO + Security approval required")
[13] evaluate_consistency()
[14] formulate_answer("No general retention beyond 90 days; extension allowed only
under approved investigation workflow")
[15] return_answer()

```

A simplified trace for Agent B may be represented as follows:

```

[1] receive_task("EU data-retention policy for failed-login IP addresses")
[2] search_policy_index("failed-login IP retention")
[3] retrieve_document("Security Logging Standard v3.2")
[4] inspect_clause("retain failed-login metadata for 90 days")
[5] search_policy_index("retention exception security investigation")
[6] retrieve_document("Incident Response Exception Policy")
[7] inspect_clause("security investigations may extend retention")
[8] generate_answer("Yes, retention may exceed 90 days for security reasons")
[9] return_answer()

```

At the surface workflow level, the two agents look similar. Both receive the task, retrieve relevant documents, inspect policy clauses, and answer. In an ordinary agentic description, both might simply be described as retrieval-enabled policy agents.

However, the agemic reduction reveals a structural difference. Agent A may be transcribed as:

$$G \rightarrow P \rightarrow R \rightarrow M \rightarrow R \rightarrow M \rightarrow F \rightarrow R \rightarrow M \rightarrow O$$

or, in compressed form,

$$G \rightarrow P \rightarrow R \rightarrow M \rightarrow F \rightarrow O$$

where **G** denotes goal commitment, **P** plan formation, **R** retrieval, **M** monitoring, **F** failure repair, and **O** closure.

By contrast, Agent B may be transcribed as:

$$G \rightarrow R \rightarrow O$$

or at most,

$$G \rightarrow R \rightarrow E \rightarrow O$$

depending on how one treats answer generation. The crucial point is that Agent B does not instantiate a genuine monitor-repair pair. It retrieves and answers, but it does not

detect that a decisive approval condition remains unchecked. Its answer is therefore surface-plausible but structurally under-controlled.

This example yields a concrete diagnostic gain. An ordinary workflow description might say only that Agent B missed an exception detail or performed insufficient retrieval. Agemic analysis permits a more precise diagnosis: retrieval occurred without stable monitoring, and answer generation was not governed by a genuine repair cycle after inconsistency detection. Agent A and Agent B are therefore *agentially similar* but *agemically distinct*.

The example also clarifies agemic equivalence. If a third agent used vector retrieval instead of keyword search but still instantiated the same sequence of goal commitment, planning, retrieval, monitoring, repair, and closure, then it would be agentially different from Agent A while remaining agemically equivalent to it.

This example shows that agemic analysis does not merely rename agent steps, but distinguishes structurally monitored retrieval-and-repair from superficially similar but under-controlled retrieval-and-answer behavior.

Agentic and Agemic Equivalence

One of the most useful consequences of the distinction is the possibility of distinguishing between similarity at the level of realized behavior and similarity at the level of agency structure.

The synthetic case above makes this distinction clearer. Agent A and Agent B are superficially similar at the agentic level. Both receive the same task, query policy sources, inspect clauses, and return an answer. If described only at the level of visible workflow, both may appear to be retrieval-enabled policy agents carrying out a multi-step compliance query.

Yet the agemic analysis shows that they are not structurally equivalent. Agent A instantiates goal commitment, explicit plan formation, retrieval, monitoring, repair, renewed retrieval, and closure. Agent B, by contrast, instantiates retrieval and answer generation but does not instantiate a genuine monitor-repair pair. It reaches an answer without structurally recognizing that a decisive condition remains unchecked. The two systems are therefore *agentially similar* in appearance but *agemically distinct* in organization.

This distinction matters because agentic resemblance can easily conceal structural difference. Two systems may both appear to be planning or retrieval agents, yet only one may realize the agemes necessary for controlled revision and recovery. In such cases, surface similarity is real but analytically insufficient.

Conversely, two systems may be *agentially different* while remaining *agemically equivalent*. Suppose Agent C performs the same task as Agent A but uses vector

retrieval rather than keyword search, invokes a different summarization tool, and structures its intermediate reasoning differently. At the level of concrete realization, Agent C differs from Agent A in tools, trace details, and execution style. Yet if Agent C still instantiates the same structural sequence of goal commitment, planning, retrieval, monitoring, repair, and closure, then Agent C is agentially different from Agent A while remaining agemically equivalent to it.

Agemic equivalence is therefore not behavioral identity. It is the equivalence of a structurally significant agency organization after abstraction from incidental realization. Agent A and Agent C may be agemically equivalent despite different concrete mechanisms, whereas Agent A and Agent B may be agentially similar yet agemically distinct because the latter lacks a genuine monitor-repair structure.

The gain of this distinction is analytic and diagnostic. It permits comparison of systems not only by what they visibly do, but by what units of agency they truly realize.

Relation to Concrete and General Levels in Pattern-Oriented Diagnostics

This distinction aligns naturally with the broader Pattern-Oriented Diagnostics contrast between the concrete and the general⁴.

Agentic AI belongs more naturally to the concrete side. It concerns actual systems, concrete workflows, specific traces, explicit tool calls, observed plans, realized behaviors, and manifested artifacts.

Agemic AI belongs more naturally to the general side. It concerns invariant structures, functionally significant distinctions, abstract units of agency, recurring compositions, and patterns that hold across many implementations.

This does not mean that Agentic AI is exhausted by the concrete or that Agemic AI is detached from realization. Rather, the distinction indicates a direction of abstraction.

Agentic AI studies agents as manifested wholes.

Agemic AI studies the pattern language from which agency is composed.

This makes Agemic AI especially suitable for formal theory, comparative analysis, and higher-order diagnostics.

⁴ Pattern! What Pattern? Theoretical Software Diagnostics, Fourth Edition, page 216 (<https://dumpanalysis.org/pattern-square>)

A Pattern-Oriented Diagnostics Interpretation

From the viewpoint of Pattern-Oriented Diagnostics, the value of Agemic AI lies not merely in naming parts of an agent, but in reducing observed behavior to structurally interpretable units. Agentic systems emit traces, logs, plans, memory accesses, tool calls, retries, critiques, and other execution artifacts. These describe what happened, but they do not always reveal which units of agency were active, absent, unstable, or miscomposed.

Agemic analysis is intended to supply that missing level. A surface description may say that a system searched, executed, retried, and eventually failed. An agemic description may instead reveal that retrieval occurred without stable monitoring, that repeated execution was not guided by a genuine repair cycle, or that delegation displaced rather than preserved goal commitment. In this way, agemic analysis seeks to transform surface sequences of events into a smaller structural account of agency organization.

The diagnostic gain lies precisely here. An ordinary workflow description may tell us that a system retried several times. Agemic analysis asks whether those retries instantiated repair in the stronger sense, that is, whether they were governed by monitoring, evaluative discrimination, and structurally meaningful revision. A superficially similar trace may therefore correspond either to genuine repair or to blind repetition. The distinction is behavioral at the surface, but structural at the agemic level.

Agemic Diagnostics of Failure

The diagnostic value of Agemic AI becomes clearest in failure analysis. A system may fail not because it lacks intelligence in some vague global sense, but because a particular unit of agency is absent, unstable, or miscomposed.

A model may answer incorrectly because retrieval occurred without adequate monitoring. A system may loop indefinitely because closure is weak, and self-monitoring cannot distinguish between progress and repetition. A multi-agent architecture may collapse into confusion when delegation is present, but coordination is absent, or when delegation overwhelms stable goal commitment. In each case, the point is not merely that the system behaved badly, but that a particular organization of agency failed to form or failed to persist.

Agemic analysis, therefore, aims to decompose failure into structurally interpretable deficits. It replaces undifferentiated diagnoses such as bad planning, hallucination, or poor autonomy with a more explicit account of which units of agency were missing, weak, excessive, or in conflict. To the extent that this succeeds, it offers a more precise basis for both interpretation and redesign.

Agemic Repertoire and Minimal Agency

The notion of an agemic repertoire also allows one to ask what counts as minimal agency.

Not every intelligent system is agentic. Not every sequence of outputs manifests a stable purposive structure. Agemic analysis provides a way to state why. A system may possess some agemes but not others. It may retrieve but not plan. It may select tools but not monitor results. It may maintain memory but not revise goals. It may terminate outputs but not have closure in the stronger sense of resolved task completion.

Minimal agency may therefore be understood not as a binary property but as the realization of a minimal agemic repertoire. One may imagine threshold repertoires for different classes of systems. A simple assistant may require goal commitment, execution, and closure. A more robust autonomous system may require planning, monitoring, repair, and memory updates as well. A collaborative multi-agent system may also require delegation and coordination.

Such thinking supports a more refined taxonomy of artificial agency than the broad and often overextended label “agentic.”

Phonetic Variation, Allophony, and Surface Realization

The phonetics and phonemics analogy can be developed further.

In linguistics, a single phoneme may have multiple phonetic realizations depending on context. These variants are often called allophones⁵. By analogy, one may say that a single ageme can admit multiple agentic realizations. The repair ageme, for example, might be realized as explicit self-critique in one system, as evaluator feedback in another, or as constraint-based re-execution in a third. The underlying agemic role may remain stable even though its manifestation differs.

This point strengthens the analogy. It shows that Agemic AI does not require a rigid one-to-one correspondence between agemes and implementation modules. The same ageme may be realized differently in different architectures, just as the same phoneme may be realized differently in different phonetic contexts.

At the same time, small differences in realization may or may not matter agemically. Some implementation differences are merely surface variations. Others are structurally decisive. Determining which is which is precisely part of the task of Agemics.

⁵ <https://en.wikipedia.org/wiki/Allophone>

Surface Agency and Deep Agency Structure

One may therefore distinguish between surface agency and deep agency structure.

Surface agency corresponds to the visible, logged, or externally inspectable sequence of agent-like behavior. It includes the steps that are visible in traces and workflows.

Deep agency structure corresponds to the agemic organization that makes those surface steps intelligible as a structured form of agency. It includes the repertoire of agemes, the oppositions among them, and the ways they are composed.

This distinction is useful because many current systems are evaluated primarily at the level of surface agency. They appear agentic because they produce plans, invoke tools, or loop through subtasks. But the deeper question is whether these observed behaviors correspond to a stable, interpretable agency structure. Agemic AI provides the framework for that deeper question.

Illustrative Comparison

The distinction may be summarized concisely. Agentic AI concerns the realized behavior of systems that act as agents. Agemic AI concerns the minimal functionally contrastive units from which such behavior is structurally composed. The former describes agentic manifestation. The latter seeks a reduced structural account of agency organization. The concepts are therefore complementary rather than competitive. Agemic AI does not replace Agentic AI. It introduces a second level of analysis.

Originality and Conceptual Novelty

The originality of Agemic AI does not lie in the broad claim that agency has internal structure. That claim already belongs to established traditions of agent theory, cognitive architecture, and contemporary agentic system design. Nor does the present proposal claim that planning, monitoring, retrieval, delegation, memory, or execution have never been distinguished before.

Its novelty lies elsewhere.

First, it introduces a distinct analytical level between whole-agent description and implementation detail. This level is defined not by modules, stages, or capabilities as such, but by *minimal functionally contrastive units of agency*.

Second, it makes the analogy with phonetics and phonemics methodologically central rather than merely rhetorical. The proposal is that realized agentic behavior admits a

structurally reduced description in terms of contrastive agency units, just as realized speech admits a structurally reduced description in terms of contrastive sound units.

Third, it proposes explicit criteria for individuating such units. Agemes are not merely named because they seem intuitively important. They are proposed as candidates only insofar as they satisfy criteria of contrast, realizational plurality, minimality, compositionality, and diagnostic yield.

Fourth, it connects this structural level to diagnosis. The proposal is not merely to redescribe agent architectures, but to use agemic analysis to distinguish structural presence from absence, genuine repair from blind repetition, and surface similarity from deeper agency equivalence.

The originality of Agemic AI, therefore, lies not in the generic decomposition of agency but in the introduction of a criterion-guided, phonemics-inspired structural level for the analysis and diagnosis of artificial agency.

Within Pattern-Oriented Diagnostics, earlier trace and log analysis patterns already anticipated important elements of Agemic AI. Opposition Messages⁶ introduced structurally meaningful binary contrasts such as open/close, allocate/free, and call/return, showing that diagnostic meaning often arises not from isolated messages but from their opposition, symmetry, and incompleteness. Activity Theatre⁷ reduced traces to scripts of activity roles, while Motifs⁸ identified recurrent higher-order trace structures. Agemic AI extends this line of thought by asking whether such contrastive and scriptable structures can be grounded in minimal functionally contrastive units of agency, agemes. In this sense, Opposition Messages may be seen as a precursor of agemic opposition, Activity Theatre as a precursor of agemic transcription, and Motifs as recurrent compositions of agemes rather than agemes themselves.

Implications for Future Work

Once the distinction is introduced, many lines of development become possible.

One may construct typologies of agemes for different classes of systems. One may describe canonical agemic patterns for retrieval agents, planning agents, coding agents, evaluator agents, and collaborative agent societies. One may define agemic deficiencies and pathologies. One may compare systems not only by benchmark performance but by agemic repertoire and agemic organization. One may design logging and tracing infrastructures that expose agentic behavior in ways suitable for agemic transcription and diagnosis.

⁶ Opposition Messages. Trace, Log, Text, Narrative, Data, Fifth Edition, page 225

⁷ Activity Theatre. Ibid., page 42

⁸ Motif. Ibid., page 214

From the perspective of Pattern-Oriented Diagnostics, one may also imagine a repertoire of analysis patterns⁹ specifically for agemic interpretation of traces, logs, and execution records from AI systems. Such work would help transform current agent observability into a more disciplined structural diagnostics of agency.

The present framework remains preliminary. Its strongest future requirement is empirical. A mature Agemics of artificial agency will require worked examples, comparative trace analyses, and practical demonstrations that agemic description yields explanatory or diagnostic advantages over ordinary workflow description alone.

Conclusion

The distinction between Agentic AI and Agemic AI is useful only if it yields more than a new vocabulary. The argument of this article is that such a yield becomes possible once agency is analyzed not merely at the level of whole systems or implementation components, but at the level of minimal functionally contrastive units. These units, here called agemes, are proposed as the basis for a structurally reduced account of agentic behavior.

The contribution of the article is therefore twofold. Conceptually, it distinguishes between realized agency and agency structure. Methodologically, it proposes criteria for individuating agemes and a procedure for reducing concrete traces to agemic transcriptions. The framework remains preliminary, but its ambition is clear: to enable a more disciplined comparison and diagnosis of artificial agents than surface-level workflow descriptions alone can provide.

In its strongest formulation, the distinction may be stated as follows:

Agentic AI concerns realized agents. Agemic AI concerns the minimal contrastive units from which agency is structurally composed.

⁹ Pattern Repertoire. Theoretical Software Diagnostics, Fourth Edition, page 229 (<https://www.dumpanalysis.org/pattern-repertoire>)

Vocabulary

Agentic AI

Artificial intelligence considered at the level of realized agency. It concerns systems that behave as agents by receiving or formulating goals, selecting means, invoking tools, monitoring outcomes, revising behavior, and sustaining purposive action across time.

Agemic AI

The study of the minimal functionally contrastive units of agency from which agentic behavior is composed. It concerns not the realized whole agent as such, but the abstract structural units and relations that make such agency possible.

Ageme

The smallest functionally contrastive unit of agency in an AI system. It is not simply any module, event, or computational step. It is a unit whose presence, absence, replacement, or recombination significantly alters the structure of agency.

Agemics

The field concerned with the study of agemes, their compositions, their oppositions, and their recurrent structural organizations.

Agemic Transcription

A reduced representation of agentic behavior in terms of agemes rather than raw concrete events. It abstracts from implementation detail and preserves only those distinctions that matter structurally for agency.

Agemic Equivalence

Structural equivalence at the agemic level. Two systems are agemically equivalent when, despite differences in architecture, prompting, tools, code, or execution traces, they instantiate the same or isomorphic agemic structure.

Agemic Deficiency

The absence, weakness, or instability of an ageme required for a certain class of agency. A system may be deficient in monitoring, repair, closure, delegation, or another agency unit, and such a deficiency may explain observed failure.

Agemic Analysis

The structural interpretation of agentic behavior in terms of agemes, agemic repertoires, agemic patterns, agemic deficiencies, and agemic equivalences. Within Pattern-Oriented Diagnostics, this term may be used for the analysis that moves from observable AI artifacts to the structure of agency they instantiate.