

117-255-D Cognitome: The Changing Boundaries of the Knowledge System for an Expert

When Expertise Stops Residing Inside a Single Human Mind

Introduction — A Quiet Change in the Meaning of Expertise

For centuries, the expert was relatively easy to identify.

- A mathematician knew mathematics.
A physician knew medicine.
An attorney knew law.
An engineer knew engineering.
- Tools certainly assisted them—books, calculators, databases, laboratories, computers—but the recognized locus of cognition remained the human expert.
- The traditional knowledge-production model could therefore be represented approximately as:

Expert → Reasoning → KnowledgeArtifact

- Artificial intelligence is beginning to disturb this deceptively simple architecture.
- An expert can now formulate a hypothesis with an AI model, ask another model to challenge it, retrieve literature through machine assistance, generate simulations, have an agent write code, ask another AI to analyze the results, solicit human peer review, revise the argument with AI assistance and eventually publish an artifact whose intellectual genealogy may contain contributions from numerous human and machine Actants.
- The important transformation may therefore not be:
- AI replaces the expert.
- It may be:
- The boundary of the expert's knowledge system is changing.
- That is a Cognitome problem.

JUST IN: OpenAI pulls sponsorship of a Caltech math hackathon after mathematicians revolt

- Academics call it "slop mathematics," credit taken, verification uncompensated
- Came right after OpenAI's Navier-Stokes claim, possibly built on researchers' own chats

Anthropic hasn't said if it's staying in.



BUSINESS INSIDER

OpenAI pulls out of Caltech math hackathon after mathematicians issue open letter criticizing AI math 'slop'

1. The Classical Expert

- Consider a mathematician.
- The classical architecture might be simplified as:
Problem → Mathematician → Reasoning → Proof
- Of course, the mathematician used external resources:
Books, Papers
, Colleagues
, Calculators
- But these generally occupied recognizable supporting roles.

- The intellectual boundary remained relatively clear.
- The mathematician was the principal cognitive engine.
- The proof was attributed to the mathematician.
- Other mathematicians independently verified it.
- The institution published it.
- Knowledge had a reasonably understandable chain of cognitive custody.

2. Enter the AI Collaborator

- Now imagine the same mathematician working in 2026.
 - The mathematician proposes an idea.
 - An LLM explores it.
 - The mathematician rejects three approaches.
 - The model generates a fourth.
 - The mathematician recognizes something promising in it.
 - Another model searches for counterexamples.
 - An agent writes numerical experiments.
 - A theorem prover verifies part of the argument.
 - Another human reviews the result.
 - The original mathematician modifies the conjecture.
 - The AI produces an alternative derivation.
 - Eventually a publishable result emerges.
- What is the cognitive system now?
- Certainly not simply:

$\text{Mathematician} \rightarrow \text{Proof}$
- A better representation might be:

$\text{Human} \leftrightarrow \text{LLM} \leftrightarrow \text{Tools} \leftrightarrow \text{Literature} \leftrightarrow \text{OtherHumans} \leftrightarrow \text{Verification}$
- The knowledge-producing system has expanded.

3. The Expert Has Acquired a Cognitome

- This is where the GUDIYA concept of the Cognitome becomes useful.
- The meaningful unit of cognition is increasingly not:

Human

 nor merely:

$\text{Human} + \text{AI}$

 but a dynamically assembled cognitive ecosystem:

$\text{ExpertCognitome}(t) = \{\text{Human, Models, Agents, Tools, Memory, Data, Literature, Humans, Institutions}\}$

 connected through changing relationships.
- The expert remains extraordinarily important.
- But the expert is increasingly operating as the organizing intelligence within a larger cognitive system.

4. The Boundary Is Becoming Porous

- This creates an intriguing boundary problem.
- Suppose a mathematician asks an AI:
 - "Does this argument hold?"
 - The AI identifies a missing assumption.

- The mathematician fixes it.
- Who contributed to the proof?
- Now suppose the AI proposes the missing lemma.
- The mathematician proves the lemma.
- Now suppose another AI proves it.
- Then a formal verification system checks it.
- Where exactly does the mathematician's cognition end?
- Where does machine cognition begin?
- And perhaps more importantly:
 - Is that even the correct system boundary anymore?
- The traditional boundary surrounded the individual: [Human Expert]
- The emerging boundary may surround the cognitive ensemble:

[Human+AI+Tools+KnowledgeSources+Verification]
- The unit of expertise is beginning to expand.

5. This Is Not Simply Intelligence Amplification

- Humans have always used tools.
- A calculator amplifies arithmetic.
- A telescope amplifies vision.
- A database amplifies retrieval.
- But generative and agentic AI introduce something qualitatively interesting.
- The tool can now participate in:
 - interpretation,
 - hypothesis generation,
 - critique,
 - synthesis,
 - planning,
 - experimentation,
 - coding,
 - verification,
 - reformulation.
- The tool therefore enters portions of the loop previously classified as reasoning itself.
- That makes the boundary much harder to draw.

6. Knowledge Production Becomes Multi-Actant

The Cognitome therefore suggests another representation.

Traditional:

$$K = f(H)$$

where H represents the human expert.

Emerging:

$$K = f(H, A_1, A_2, \dots, A_n, D, T, E)$$

where:

- H= human expert,
- A_i = AI Actants,
- D= data and accumulated state,
- T= tools,
- E= external knowledge environment.

Knowledge becomes the output of a coupled cognitive system.

This does not imply that all participants contributed equally.

Quite the opposite.

Determining who contributed what becomes increasingly important.

7. The Expert's Most Valuable Skill May Move Upward

- This suggests an interesting evolution of expertise.
- When calculation was scarce, calculation mattered enormously.
- When information was scarce, possessing information mattered enormously.
- When analysis was scarce, analytical ability became differentiating.
- AI makes many of these capabilities abundant.
- The expert's comparative advantage may therefore migrate upward toward:
 - ProblemFraming
 - Judgment
 - Interpretation
 - EpistemicDiscipline
 - SystemArchitecture
 - Verification
 - Responsibility
- The future expert may increasingly resemble the chief engineer of a cognitive system.
- The mathematician does not disappear.
- The mathematician becomes capable of orchestrating vastly more mathematical cognition.

8. Expertise Becomes Orchestration

- Consider two equally accomplished scientists.
- Scientist A works alone with conventional tools.
- Scientist B orchestrates:
 - several frontier models,
 - specialized agents,
 - simulation tools,
 - literature retrieval,
 - formal verification,
 - human collaborators.
- Their underlying biological intelligence may be comparable.
- Their effective cognitive capacity may not be.
- We could therefore begin distinguishing:

IndividualExpertise

from:

OrchestratedExpertise

- The latter is a property of the expert's Cognitome.

9. This Changes the Meaning of Intellectual Productivity

- The traditional model often implicitly measured:
 - Productivity $\approx$ IndividualCognitiveOutput
- The emerging model may look more like:
 - Productivity = f(Expertise, CognitomeArchitecture, AILeverage, Verification, Coordination)
- The best researcher may eventually not be the person who personally performs every intellectual operation.

- It may be the person capable of constructing and directing the most effective knowledge-producing Cognitome.
- This resembles other engineering transitions.
 - The best aircraft designer does not personally machine every turbine blade.
 - The chief engineer understands how the entire system produces the desired outcome.
 - Knowledge work may be undergoing a similar transition.

10. But Attribution Becomes Difficult

- The mathematics controversy that motivated this chapter points toward precisely this problem.
- When knowledge passes repeatedly between:

$$\text{Human} \leftrightarrow \text{AI}$$
 the provenance of ideas becomes difficult to reconstruct.
- Suppose:

$$H_1 \rightarrow \text{AI} \rightarrow H_2 \rightarrow \text{AI} \rightarrow H_3$$
 and H_3 eventually produces an important result.
- The relevant question becomes:
- Where did the decisive cognitive contribution originate?
- Traditional citation mechanisms were designed primarily for:

$$\text{Human} \rightarrow \text{Human}$$
- They were not designed for persistent, recursive human-machine cognitive interaction.

11. Knowledge May Need Cognitive Provenance

- The Cognitome therefore suggests a future requirement:
Cognitive Provenance
 - A knowledge artifact may need an auditable history describing:
 - originating hypotheses,
 - human contributors,
 - AI contributors,
 - source material,
 - model interactions,
 - transformations,
 - verification steps,
 - rejected branches,
 - final human judgments.
 - Conceptually:

$$\text{Idea} \rightarrow \text{Transformation} \rightarrow \text{Inference} \rightarrow \text{Verification} \rightarrow \text{Artifact}$$
 becomes a traceable lineage.
 - This is not merely an academic-credit problem.
 - It is an epistemic problem.

12. Chain of Cognitive Custody

- Borrowing from forensic practice, we might call this:

Chain of Cognitive Custody
- For consequential knowledge, we increasingly need to know:
- Where did this claim come from?
- Which Actant transformed it?
- What evidence entered the reasoning?
- Which model produced which inference?

- What did the human accept or reject?
- Who verified the result?
- Which state persisted between interactions?
- The final artifact alone may no longer tell us enough about how the knowledge came into existence.

13. Verification Becomes More Important, Not Less

- AI can dramatically lower the cost of producing candidate knowledge.
- But lowering production cost does not necessarily lower verification cost.
- Indeed:

can produce: $\text{Cost}_{\text{Generation}} \downarrow$

which may cause: $\text{Volume}_{\text{Claims}} \uparrow$

$\text{VerificationBurden} \uparrow$

- This explains some of the tension now appearing in scientific communities.
- The issue need not be interpreted simply as academics resisting AI.
- There is a legitimate systems question:
 - If machines dramatically increase the rate at which candidate knowledge enters the system, who supplies the corresponding verification capacity?
- That is a stability problem.

14. Cognitive Velocity Appears Again

- The GUDIYA Cognitome already gives us language for this.
- Suppose AI-assisted conjecture generation operates at:

V_G

while rigorous human verification operates at:

V_V

- If:

$$V_G \gg V_V$$

the knowledge system accumulates a growing verification deficit.

- That resembles a Cognitive Velocity mismatch.
- The system can become saturated with plausible but insufficiently verified claims.
- Thus the AI mathematics controversy potentially points toward something deeper than professional insecurity.
- It may expose a new Rendezvous problem between machine-speed generation and human-speed verification.

15. Expertise Is Becoming a Dynamic Boundary

There is an even deeper connection to our definition of the Cognitive Field as an open system.

- At 9:00 AM, an expert's working Cognitome might contain:

$\{H, A_1, \text{Database}\}$

- At 10:00:

$\{H, A_1, A_2, \text{Database}, \text{Agent}_1\}$

- At noon:

$\{H, A_2, \text{Agent}_1, \text{Colleague}\}$

- At 3:00:

{H, A₃, FormalVerifier, Colleague, Dataset}

- Thus:

$$\text{ExpertCognitome}(t + \Delta t) \neq \text{ExpertCognitome}(t)$$

- The expert's effective knowledge-production system is an open cognitive system.
- Its membership changes while the work is occurring.

16. AI Unbounds the Expert

- This connects directly with GUDIYA's recent observation that Agentic AI unbounds Cognitive Space.
- Previously, the practical boundary of an expert was constrained by:
 - available time,
 - memory,
 - accessible literature,
 - personal skills,
 - available collaborators.
- Agentic AI weakens several of those constraints.
- An expert can suddenly reach into:
 - unfamiliar disciplines,
 - enormous literatures,
 - programming languages,
 - simulation environments,
 - datasets,
 - specialized models,
 - autonomous research workflows.
- The expert has not become omniscient.
- But the reachable cognitive state space has expanded dramatically.

17. The Expert Becomes a Cognitive Enterprise

This suggests a useful conceptual leap.

Historically:

Expert $\approx$ Individual

Increasingly:

Expert $\approx$ Micro Cognitive Enterprise

- One human may command:
 - research agents,
 - coding agents,
 - critic agents,
 - search agents,
 - simulation agents,
 - verification systems,
 - persistent memory.
- A single scientist may eventually operate a cognitive organization that previously would have required a laboratory.
- A single engineer may operate something resembling a small engineering department.

- A single mathematician may orchestrate something resembling a mathematical institute.
- This is a profound democratization of cognitive capacity.

18. But Larger Cognitomes Create New Stability Problems

- Amplification creates complexity.
- As the expert's Cognitome grows:
$$N_{\text{Actants}} \uparrow$$
potential Rendezvous relationships can grow dramatically.
- Then appear:
 - contradictory outputs,
 - recursive loops,
 - stale context,
 - persistent objectives,
 - attribution ambiguity,
 - velocity mismatches,
 - accumulated cognitive state,
 - verification deficits,
 - coordination failures.
- The expert may become vastly more capable while simultaneously operating a more complex system.
- This is precisely the territory of Stability Engineering.

19. The Question Changes

- The conventional AI question asks:
 - Will AI replace the mathematician?
- The Cognitome asks something more interesting:
 - What is a mathematician when mathematical cognition is distributed across a dynamically changing ensemble of humans, models, agents, tools and verification systems?
- Likewise:
 - What is a physician?
 - What is an attorney?
 - What is an engineer?
 - What is a scientist?
 - The professions may remain.
 - Their cognitive boundaries may not.

20. A New Definition of the Expert

We can therefore propose:

Cognitome Definition of an AI-Era Expert

An expert is increasingly not merely an individual possessing specialized knowledge, but the responsible human authority at the center of a dynamically assembled knowledge system composed of human judgment, machine cognition, tools, memory, evidence, verification mechanisms and other cognitive Actants.

- The crucial phrase is: responsible human authority.
- Cognitive work can distribute.
- Responsibility cannot simply evaporate into the network.

21. The Expert's New Instrument Panel

- If this evolution continues, experts may eventually require something analogous to an aircraft cockpit display.
- Not simply:
- What did the AI say?
- But:
 - Which Actants contributed?
 - What evidence was used?
 - What remains unverified?
 - Where did disagreement occur?
 - What is the provenance of this conclusion?
 - How much machine-generated reasoning accumulated?
 - What assumptions propagated across agents?
 - Which results are independently corroborated?
 - Where are the uncertainty concentrations?
- In GUDIYA language, the expert may eventually need a GCR for knowledge production.
- The expert must see the state of the Cognitome.

22. From Personal Computer to Personal Cognitome

- There may even be a historical symmetry here.
- The computer revolution gave the expert:

Personal Computer
- The AI revolution may give the expert:

Personal Cognitome
- The first amplified computation.
- The second amplifies cognition itself.
- And just as personal computers eventually required operating systems, networks, cybersecurity and governance, Personal Cognitomes may eventually require:
 - identity,
 - provenance,
 - orchestration,
 - observability,
 - verification,
 - stability engineering.

23. The Deeper Lesson from the Mathematicians

- The current unease in mathematics should therefore not be reduced to:
- “Mathematicians are afraid AI will take their jobs.”
- That would miss the more interesting transformation.
- The disturbance is occurring at the boundary of the knowledge system itself.
- Questions of:
 - authorship,
 - credit,
 - verification,
 - provenance,
 - contribution,

- intellectual ownership
- become difficult precisely because the old boundary: [Expert]
is giving way to:
[Expert+AI+Agents+Tools+Knowledge+Verification]
- The institutional rules were built for the first architecture.
- Knowledge production is beginning to resemble the second.

24. Canonical Cognitome Principle

The Expanding Expert Boundary Principle

As AI participates increasingly in reasoning rather than merely information retrieval or computation, the effective boundary of an expert's knowledge-producing system expands beyond the individual. Expertise becomes a property of a dynamically assembled Cognitome containing humans, models, agents, tools, memory, evidence and verification mechanisms.

And its corollary:

The central AI-era question is therefore not simply whether machines replace experts, but how expertise, responsibility, provenance and verification are defined when the expert becomes the orchestrator of an open cognitive system.

Conclusion — Where Does the Expert End?

- The mathematicians may only be among the first professions to confront the question clearly.
 - Medicine will confront it.
 - Engineering will confront it.
 - Law will confront it.
 - Science will confront it.
 - Management will confront it.
- Eventually society may discover that the defining transformation of professional AI was not simply that machines became intelligent.
- It was that the boundary around human expertise stopped being obvious.
- The expert did not disappear.
- The expert's cognitive perimeter expanded.

Expert → Expert + AI → ExpertCognitome
- And perhaps the most important question for the coming knowledge economy will no longer be:
 - How intelligent is this expert?
- It may become:
 - How capable, trustworthy, verifiable—and stable—is the cognitive system this expert commands?
- That is where the Cognitome enters the future of expertise.