

Leadership Track
Beyond LLMs:
Cognitive AI, and the Future of AI

Peter Voss

Summer AI Intensive
August 2026



Learning Outcome

Strategic Overview of AI

Generative AI is unlikely to scale to AGI. Cognitive AI promises to overcome current limitation.

- What makes human intelligence so special
- Strengths and weaknesses of Generative AI
- The importance of incremental, real-time learning
- DARPA's "Three Waves of AI"
- AI from first principles: Cognitive AI
- A roadmap to AGI

What makes Human Intelligence so Special?

- ✓ Adaptive
- ✓ Conceptual
- ✓ Metacognitive
- ✓ Tool User

What makes Generative AI so Powerful?

- ✓ Massive Data and Compute
- ✓ Natural Language Interface
- ✓ APIs
- ✓ Inference Speed
- ✓ Specialized Versions

Inherent LLM Limitations

Cannot Learn Incrementally

Hallucinations, Reliability Issues

Not Auditable, Needs Human in the Loop

Privacy and Security Issues

Insane Cost, Diminishing Returns

Environmental Impact

What Generative AI is Good at

- ✓ Search, and helping with Research
- ✓ Prototyping and standard Coding Tasks
- ✓ Image generation
- ✓ Non-critical Agent Tasks
- ✓ Specialized Narrow Applications

What Generative AI is *not* Good at

- ✓ Critical, private, or secure applications
- ✓ Apps without a human in-the-loop
- ✓ Apps requiring real-time adaptation
- ✓ Apps requiring robust higher-level cognition
- ✓ Data- and power-efficient apps

The Wall is Real!

Computer Science > Artificial Intelligence

[Submitted on 25 Jul 2025 (v1), last revised 30 Jul 2025 (this version, v2)]

The wall confronting large language models

Peter V. Coveney, Sauro Succi

We show that the scaling laws which determine the performance of large language models (LLMs) severely limit their ability to improve the uncertainty of their predictions. As a result, raising their reliability to meet the standards of scientific inquiry is intractable by any reasonable measure. We argue that the very mechanism which fuels much of the learning power of LLMs, namely the ability to generate non-Gaussian output distributions from Gaussian input ones, might well be at the roots of their propensity to produce error pileup, ensuing information catastrophes and degenerative AI behaviour. This tension between learning and accuracy is a likely candidate mechanism underlying the observed low values of the scaling components. It is substantially compounded by the deluge of spurious correlations pointed out by Calude and Longo which rapidly increase in any data set merely as a function of its size, regardless of its nature. The fact that a degenerative AI pathway is a very probable feature of the LLM landscape does not mean that it must inevitably arise in all future AI research. Its avoidance, which we also discuss in this paper, necessitates putting a much higher premium on insight and understanding of the structural characteristics of the problems being investigated.

THE ECONOMIC TIMES | Panache

English Edition ▾ | 01 September, 2025, 11:15 PM IST | Today's ePaper

NVIDIA says 'Small Language Models' are the future of AI, but why isn't anyone switching yet?
The \$57 billion reason may surprise you

TECH

OpenAI's Sam Altman sees AI bubble forming as industry spending surges

PUBLISHED MON, AUG 18 2025-3:51 AM EDT UPDATED MON, AUG 18 2025-2:24 PM EDT



Dylan Butts
@IN/DYLAN-B-7A451A107

WATCH LIVE

KEY POINTS

- OpenAI CEO Sam Altman has reportedly said that he believes AI could be in a bubble, comparing market conditions to those of the dotcom boom in the 1990s.
- “Are we in a phase where investors as a whole are overexcited about AI? My opinion is yes. Is AI the most important thing to happen in a very long time? My opinion is also yes,” he's quoted as saying.



Denis O. • 1st
Fintech Professional | AI Solution Architect | Real Time Data, Ontologies & Kn...
1yr • Edited •


RAG is nothing but a bandaid to the underlying LLM issue (or feature rather) that cannot be solved without significant advances. Solving hallucinations may require new AI/computing paradigms. Not merely adjustments in LLM.

"Eliminating hallucinations through RAG is impossible."




RAG Does Not Reduce Hallucinations in LLMs—Math
Deep Dive
medium.com

External 'Guardrails' and Memory



- ❑ Tail wagging dog
- ❑ Performance Issues
- ❑ Cost Issues
- ❑ Does not solve multi-step reliability ($95\% ^ 8 = 66\%$)

Modern Neuro-Symbolic Harnesses



Can we have the Best of Both Worlds?

- ✓ Adaptive, reliable, and autonomous
- ✓ High-level cognition, and common sense
- ✓ Efficient, knowledgeable, and fast
- ✓ Private and secure

This is called

AGI

“The ability to learn any human cognitive task
with no more data or time than a human”

For example:

- Onboard a novel remote task as easily as a human
- Learn to drive a car in a few dozen hours

Real-Time Learning and Memory is Critical for AGI



AUG 13, 2025

“The dream of AGI continues to elude us due to its inability to continually learn.”

SAM ALTMAN
CEO, OPENAI



JANUARY 26, 2025

“Better memory, long-term memory — systems should remember the details they've been told to achieve AGI.”

DEMIS HASSABIS
CEO, GOOGLE DEEPMIND



JUNE 03, 2025

“I like to talk explain it as LLMs are a bit like a coworker with Anterograde amnesia - they don't consolidate or build long-running knowledge or expertise once training is over and all they have is short-term memory (context window).”

ANDREJ KARPATY
CO-FOUNDER, OPENAI

Real-Time Learning and Memory is Critical for AGI



FEB 11, 2025

“Current LLMs lack persistent or continuous memory, reasoning, and planning.”

YANN LECUN
CHIEF AI, META



JUNE 17, 2025

“Experience — interacting with the world — is how both people and animals learn.”

RICHARD SUTTON
TURING AWARD 2024



APRIL 04, 2025

“Does it remember not just my name, but how to pronounce my name? And when I correct it, does it remember that correction? That’s a really hard problem.”

MUSTAFA SULEYMAN
CEO, MICROSOFT AI

Real-Time Adaptive Learning

- **Incremental** — Absorbing information one piece at a time
- **Real-time** — Learning while busy with everyday activities
- **One-shot** — The ability to ‘update the model’ with single examples
- **Pro-active** — Knowing when to ask for clarification or missing details
- **Integrative** — Automatically updating and generalizing implications
- **Multi-Modal** — Learns via perception-action as well as natural language
- **Limited resources** — Learning with incomplete, noisy data in real-time
- **Autonomous** — The AI needs metacognitive control to know what it knows and what it doesn’t know, and to proactively manage the learning process

Advantages of Real-Time Incremental Learning

- 1) One million times **less training data**
- 2) Can control **data quality**, as quantity becomes a non-issue
- 3) Many orders of magnitude **less compute** needed
- 4) Can run on **local** off-the-shelf device
- 5) Allows for total control of AI, and **privacy**
- 6) Incremental learning uses **metacognition** (System 2)
- 7) Can learn **autonomously**, without costly human in-the-loop
- 8) Inherent ability to validate and reason (**no hallucinations**)

First Principles, Contrarian Insights

Real Intelligence requires:

- Adaptive Real-Time Integrative Proactive Lifelong Learning
- Metacognition (System 2) that manages Automatic Responses (System 1)
- We only need Millions of words (and 20W) to learn Language & Reasoning

Generative AI is ***inherently*** unable to update its Model in Real-Time

We need a fundamentally different Architecture: **Cognitive AI**

What is Cognitive AI?

It starts with the question: What makes Human Intelligence so special?

How we learn:


Adaptively. Conceptually. Incrementally. In real-time. Autonomously.

How we reason:

Conceptually. Contextually. Metacognitively.

The Three Waves of AI


First Wave – Traditional Programming



Wave 1

Hand-Crafted
Programming/ Logic
GOFAI


Second Wave – Neural Nets - Big Data



Wave 2

Statistical
Machine Learning
Generative AI

Third Wave – Cognitive Architecture



Wave 3

Adaptive. Contextual
Autonomous
Cognitive AI



Adapted from



Cognitive AI: A Different Architecture

- Can learn incrementally, in real time
- One- and few-shot learning
- Perceptual and natural language learning
- Adapts to changing environment & goals
- Scrutable, auditable – not a Black Box
- Neuro-Symbolic, predictable – not generative, probabilistic
- Training, adaption & operation on edge device
- Domain agnostic, general purpose intelligence engine

Statistical AI vs. Cognitive AI



Differences
Between
GenAI &
Cognitive AI

Factor	Generative AI	Cognitive AI
Focused on	Accessing Existing Knowledge	Acquiring & Applying Knowledge/ Skills Dynamically
Data Philosophy	Big Data/ Quantity	Small Data/ Quality
Training Time	Weeks or Months	In Real Time
Training Cost	Massive	Negligible
Train & Operate on Laptop	No	Yes
Hallucinates	Yes	No
Needs Human-in-Loop	Yes	Negligible
Explainable/ Scrutable	No	Yes
Direct Path to AGI	No	Yes

LLMs Cannot Get Us to AGI




“... on the path towards [AGI], an LLM is basically an off-ramp, a distraction, a dead end.”

Yann LeCun – Chief AI Scientist, Meta



“Deep learning.. [is] definitely not enough to solve AI, [not] by a long shot.”

Demis Hassabis – CEO, Google DeepMind



“We need another breakthrough... language models won't result in AGI.”

Sam Altman – CEO, OpenAI

Rockets vs Ladders



Generative AI



Cognitive AI


Limitations of Hybrid Neuro-Symbolic Systems

- Lack of common data representation
- Inability to dynamically shift between two systems
- The need for custom engineering
- The worst of both worlds:
 - Brittleness of Symbolic AI
 - Hallucinations of GenAI
 - Inability to update models incrementally in real-time

Introducing INSA:


Integrated **N**euro-**S**ymbolic **A**rchitecture

INSA: Integrated Neuro-Symbolic Architecture




- ✓ Super High Performance Vector Graph Encodes All Knowledge And Skills
- ✓ Functional components are deeply integrated with each other and with the underlying substrate-forming Graph DB
- ✓ Integrated Short- and Long-Term Memory, As Well As Dynamic Context
- ✓ Highly Effective One/ Few-Shot Learning Mechanisms

The Narrow AI Trap




Adaptive, Internal Intelligence: Nature/ Nurture

Old Commercial Aigo




New AGI Aigo



Focus on Curriculum: Toddler to Ph.D.

- Lives in a PC: Desktop vision, mouse, and text 'senses'
- Learning Natural Language & Reasoning from the ground up
- Common sense knowledge about the world, then STEM
- Gradually from Toddler, to Kid, to Teenager, to Graduate, to Ph.D.!
- Increasing learning autonomy: 'Hitting the books' (LLMs + reasoning)

The Roadmap to AGI



No Skill Left Behind!

Stag	Age	Language abilities	Reasoning / logic	Perception/visuospatial skills	Math concepts	Game understanding	Social cues & perspective-taking	Metacognitive capabilities
toddler	3	Uses 3-4 word sentences	Sorts objects by one feature (e.g., color)	Recognizes and names basic shapes (circle, square, triangle)	Counts small sets by rote (often to 3-5+)	Plays very simple turn-taking games	Recognizes basic emotions (happy, sad, mad)	Very early awareness of "I like / don't like this"
		Names common objects and people	Understands simple cause-effect ("If I push, it falls")	Matches identical pictures	Begins 1-to-1 matching	Follows 1 simple rule (roll and move, match colors)	Engages in parallel and emerging cooperative play	May say "I can't" when something feels hard but can't yet describe why
		Follows simple 1-step directions	Focuses on one aspect at a time	Notifies big/small and up/down positions	Notifies big vs small and obvious "more/less"	Understands "my turn/your turn" and simple win/lose	Responds to simple greetings and routines	
toddler	4	Asks basic "what/where" questions	Limited attention span					
		Speaks in longer, mostly clear sentences	Begins simple if-then thinking	Copies simple shapes (circle, cross)	Counts to around 10 or higher	Follows 2-step game rules	Plays cooperatively for short periods	Starts to say things like "I know this" vs "I don't know this"
		Tells very simple stories	Compares "more/less," "same/different"	Understands basic spatial words (in, on, under, beside)	Counts small sets accurately	Understands clear win/lose	Can explain another child's feelings in simple terms	May ask for help but not yet able to explain what kind of help is needed
kid	5	Asks many "why" questions	Notifies simple repeating patterns	Fits simple shapes into matching holes	Names basic shapes (circle, square, triangle)	Accepts simple randomness from dice or spinners with support	Begins to see sharing and turns as expected behavior	
		Understands in/on/under	Still intuitive rather than logical		Compares small quantities visually			
			Imaginative thinking					
kid	6	Uses complex sentences	Groups items by 2 features (color and size)	Copies more shapes (square, triangle, star)	Counts reliably to 20+	Handles basic board/card games (Candy Land, Snakes and Ladders, Go Fish)	Shares, waits, and negotiates brief turns	Begins to notice "I made a mistake" after it happens
		Mostly clear to unfamiliar adults	Follows 2-3 step instructions	Understands positional words (behind, next to, between)	Begins informal addition/subtraction with fingers or objects	Understands stable turn order and very simple strategy (collect matches)	Understands simple group rules (no cutting in line)	Can sometimes say what was hard ("It was too many")
		Retells simple stories with beginning and end	Reasons about short, everyday cause chains	Builds simple block towers with basic symmetry	Recognizes numerals 0-10+		Is motivated to follow classroom rules	Responds to prompts like "How could you do it next time?" with adult
kid	7	Answers most who/what/where questions	Emerging logical thinking for familiar situations		Compares lengths and sizes more systematically			
		Answers most why questions	Counts and categorizes					
		Understands time words (yesterday/tomorrow)	Sorts					
kid	8	Describes plans and past events with more detail	Plans simple activities	Copies complex shapes (diamond, hexagon)	Adds and subtracts small numbers (often within 10)	Remembers and applies rules from session to session	More aware of others' feelings	Can answer simple questions like "What was easy/hard?"
		Starts using some familiar non-literal phrases	Plans short sequences ("first... then...")	Understands spatial relationships (left/right, above/below)	Understands numbers as fixed quantities	Tracks simple scores	Offers basic comfort	Starts to check work when reminded
		Understands jokes and riddles	Keeps a goal in mind and chooses steps to reach it	Recreates simple patterns with blocks or tiles.	Orders numbers and begins using number lines	Manages games that require choosing between a couple of options each turn	Recognizes obvious teasing or unkindness	May adjust approach after feedback ("I'll count slower this time")
kid	9	Reads simple texts	Starts logical thinking about concrete objects		Recognizes more shapes and simple symmetry		Sees that actions can hurt feelings without physical harm	
		Explains ideas with some detail	Shows early conservation understanding	Understands maps/diagrams with basic symbols	Works with 2-digit numbers	Starts planning 1-2 moves ahead in simple strategy games (checkers, basic chess ideas)	Notifies fairness and cheating	Begins to predict performance ("I think I'll be good at this")
		Retells stories in order with key events	Considers more than one feature at a time	Rotates shapes mentally to match orientations	Begins conceptual multiplication (groups of, repeated addition)	Follows rule tweaks and variants when explained	Can talk about why rules matter	With prompts, can describe simple strategies ("I counted by twos")
kid	10	Asks clarifying questions	Understands cause-and-effect more consistently	Builds with more spatial planning	Understands tens and ones		Recognizes that different people have different preferences or opinions	Recognizes when they are confused and may ask for clarification
			Can compare, classify and sequence					
teen	11	Summarizes main points from stories or short texts	Classifies and orders objects systematically	Copies detailed figures accurately	Works confidently with larger whole numbers	Understands games with hidden information and simple probability (dice, card draws)	Reads more subtle cues (tone, expression)	Can talk about strategies that worked or didn't ("Next time I shouldn't guess")
		Uses more precise vocabulary for feelings, reasons, comparisons, descriptive language	Holds 2-3 aspects in mind (size, color, number)	Understands coordinate grids (first steps)	Adds/subtracts multi-digit numbers	Manages multiple conditions or phases without adult help	Understands being left out or included	Begins to set small goals ("I'll try to beat my score")
		Follows multi-step directions	Explains simple logical relationships	Mentally manipulates 2D shapes	Learns multiplication tables and simple division		Can describe simple friendship problems and possible	More aware of when they need help
teen	12	Gives multi-step directions		Notifies symmetry and patterns in designs	Begins fractions as parts of a whole			
		Explains thinking and problem steps	Solves multi-step problems more independently	Handles 3D spatial tasks (folding nets to cubes)	More fluent with multiplication/division	Plays strategy games with multiple ways to score or win	Recognizes social roles (leaders, followers, outsiders)	Can reflect on strengths and weaknesses in specific areas ("I'm good at reading, math is hard")
		Understands some idioms and simple metaphors	Considers alternative strategies	Reads simple maps with legend	Works with simple decimals	Adjusts strategy mid-game	Anticipates some social consequences (e.g., teasing, reasons about fairness beyond simple equal shares)	Begins to monitor progress over time
teen	13	Summarizes and compares short texts or stories	Explains why a solution works in own words	Rotates and flips shapes mentally with confidence	Works with common fractions	Understands cooperative vs competitive structures		Can follow teacher prompts to plan, check, and revise work
				Identifies missing parts in patterns	Interprets simple charts and graphs			
					Handles more complex word problems with support			
teen	14	Uses relatively sophisticated sentences	Applies logical rules to new situations	Navigates maps and diagrams fluently	Works comfortably with larger numbers, fractions, and decimals	Enjoys and manages complex strategy/Euro-style games with resources, trading, and long-term planning	Understands nuanced emotions (embarrassment, guilt, jealousy)	Begins to plan approaches ("First I'll... then I'll...")
		Participates in structured discussions	Sees patterns across experiences	Understands perspective (different views of same object)	Begins ratio and simple percent ideas	Tracks several conditions and scoring methods at once	Navigates group expectations	Can self-evaluate work against criteria
		Summarizes and paraphrases information from several sources	Evaluates strategies and justifies preferences	Mentally rotates 3D objects	Solves multi-step word problems with more autonomy		Can negotiate and compromise during conflicts	Starts to choose strategies deliberately (e.g., drawing a diagram) rather than by trial and error alone
teen	15	Explains abstract ideas with concrete examples	Can think logically across multiple variables	Excels at spatial puzzles and design tasks				
		Discusses "what if" scenarios	Begins more abstract (early formal) reasoning	Masters complex spatial reasoning (3D rotations, scale models)	Uses fractions, decimals, and ratios in multi-step problems	Thinks several turns ahead in complex games	Reads complex social cues, including subtle exclusion and peer pressure	More independently sets goals ("I want to improve my test score")
		Understands more idioms, sarcasm, and irony	Imagines multiple outcomes	Interprets detailed diagrams/maps	Starts algebraic thinking (unknowns, simple equations, function-like patterns)	Weights risk vs reward	Reflects on motives	Monitors understanding ("This part doesn't make sense yet")
teen	16	Compares arguments and identifies main claims and supporting reason	Mentally tests simple hypotheses	Excels at architecture-like tasks and video game spatial navigation	Interprets more complex graphs and data	Handles cooperative games that require role coordination and flexible planning	Is aware of social reputation and longer-term consequences	Can describe and adjust strategies without adult prompting
		Near-adult vocabulary	Evaluates strengths and weaknesses of ideas					
		Argues positions with reasons and counterarguments	Uses emerging formal reasoning more consistently	Uses advanced visuospatial skills fluidly (CAD-like mental modeling, architectural planning)	Applies pre-algebra (variables, equations, proportional reasoning)	Manages highly complex games with layered strategies and exceptions	Understands social norms and unwritten rules	Can plan, monitor, and evaluate their own learning cycle (set goal → choose strategy → check progress → adjust)
teen	17	Synthesizes information from different sources	Plans and justifies multi-step solutions	Interprets complex scientific diagrams	Works fluently with fractions, decimals, and percentages	Learns rules from written text and can teach them clearly to others	Anticipates longer-term social and emotional outcomes	Reflects on how they learn best and can articulate effective study or practice habits.
		Shifts language for audience (peers vs adults)	Reasons about abstract systems with several variables	Excels at strategy games requiring board visualization	Interprets coordinate graphs and quantitative relationships		Takes multiple perspectives and proposes balanced solutions in conflicts	

"Adapt" or "Collapse"

"Adapt" or "Collapse"

Core Intelligence, Customizable Expertise



- **Graduate Aigo** is what we offer to the world. It can be onboarded to do any deskbound job as effectively or better than human
 - **Expert Aigo** becomes expert in customer specific domains by learning using existing curricula
 - **Personalized Aigo** proactively learns about users and hyper-personalizes experiences based on user's history, preferences, goals and values
-
- ✓ Privacy and safety is guaranteed by hosting Aigo behind customer's firewall or on edge device
 - ✓ We will not sell or share data
 - ✓ Aigo is trained to boost human agency & flourishing

Commercial Advantages

- Massively lower set-up and maintenance costs
- Much better reliability, dependability, predictability
- Can ***train*** and run on Edge Device
- Will adapt to changing circumstances, learn on-the-fly
- User is not subject to the fortunes and whims of a supplier
- More intelligent via robust ontology and meta-cognition
- Scrutable. Not a black box.

We're at a Unique Moment in History...

AGI: Bigger than Electricity

Boost Human Flourishing

Tremendous Opportunities


The background of the image is a clear, vibrant blue sky filled with numerous white, puffy clouds of varying sizes. The clouds are scattered across the frame, creating a sense of depth and openness. The lighting is bright, suggesting a sunny day.

Imagine...

Boosting Human Flourishing

- ✓ **AGI Researchers in every field**
- ✓ **Massive cost reduction in goods & services**
- ✓ **Personal **Personal** Assistant to optimize individual life**

With all the Benefits of Computers...





Medical

- End Disease
- Neural & Other Implants
- Organ Replacement
- Radical Longevity/ Rejuvenation






Energy


- Batteries/ Supercapacitors
- Effective Pollution Scrubbing
- Fusion & Thermal





Nanotechnology

- Molecular Nanotech
- Nanobots (Medical, Environmental)
- Materials, Computers





Computation/ AI

- Massive Increase in Speed & Memory
- Quantum Computers
- Ongoing AI Improvements





Food

- High-Density Food Production
- Great Artificial Meat
- Optimized Healthy, Tasty Food
- Personalized, Customized Meals





Philosophy/ Psychology

- Rationality
- Rational Ethics
- Mental Health & Clarity
- Law & Order. Just Justice.
- Rational Governance

Imagine...

Using the power of billions of human-level AGI agents
to usher in the next wave of the Enlightenment –
to significantly boost **Human Flourishing**

Thank you!

peter@Aigo.ai