

IGNITION · WORKSHOP 3

AI Under the Hood

Part 2: Inside the brain

Thursday, June 25, 2026 · 1:00pm ET

Tom Bernardin

Executive Director, CDSAI · campusGenAI

OVERVIEW

What we'll cover

1 What is the AI “brain” and how is it made?

2 How does AI “think”?

3 What do we think about AI? — how our assumptions shape our work

LEARNING OUTCOMES — BY PATH

EXPLORERS

Appreciate that an AI model is trained to validate your framing — it has no independent basis for judgment beyond extending the patterns it has seen.

BUILDERS

See how different use cases are shaped by how the model works and was built — why it's strong for some tasks and unreliable for others, and what that means for where you point it.

CREATORS

Consider what it takes to augment business processes with AI at scale — the risks it introduces (drift, hallucination, bias, accountability) and how to start thinking about managing them.

FOR EVERYONE

An AI model is not analogous to a human mind — it's a pattern-predictor based on the data it was trained on. It doesn't look things up, and has no data stored anywhere.

BRIEF HISTORY

AI is not new, but GenAI is



A tool for identifying a whale from an aerial image, automatically marks its head, tail, and size estimate.



A tool for detecting extreme speech within YouTube videos.



A tool that rapidly analyzes the before and after satellite images of affected areas to assess damage.



A tool that analyzes aerial images to determine vehicle usage and density for environmental impact.

BRIEF HISTORY


“AI” vs. “Generative AI”



Deep Learning



**AI beating Humans
at GO**



**AlphaFold Protein
Folding**




Computer Vision



LLMs

HOW IT'S MADE

What shapes an AI interaction?




THE SCALE OF TRAINING

~15 trillion words (≈ tens of millions of books) · ~16,000 data-center GPUs running ~54 days (~2 months)
≈ 190,000 gaming GPUs (~tens of thousands of years on one) · ~30 GWh ≈ a year's electricity for ~3,000 homes

illustrative — comparable public model (Llama 3.1 405B); exact figures for Opus aren't published


HOW IT'S MADE

What shapes an AI interaction?



Interactions with LLMs are impacted by multiple layers

Why can it “reason” through some problems but not others?



■ its move — and the best move

A FAMILIAR OPENING

Best move: the Ruy Lopez — the classic opening move

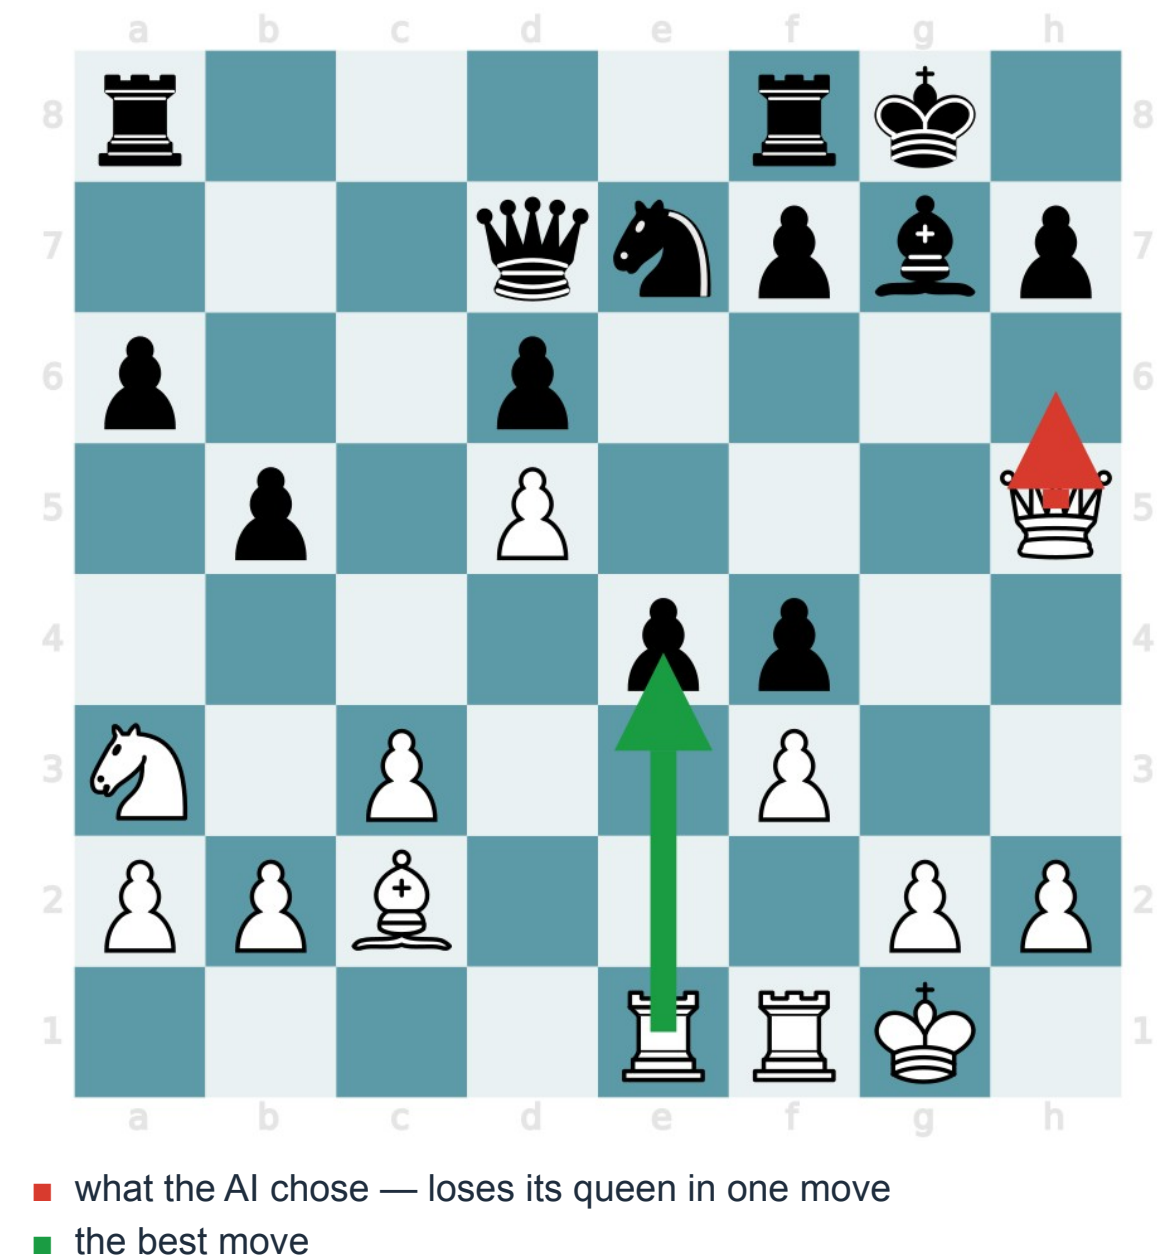
WHAT OPUS DID

Named it instantly — the move every world champion from Steinitz to Carlsen has played. Pure pattern-matching from a position it has seen countless times. And it's right.

GOOD AT — Pattern recognition

- Variants on patterns it has seen many times in its training
- Conforming your material to patterns in training data
- Generic prompt in, generic but competent answer out

Why can it “reason” through some problems but not others?



AN UNFAMILIAR MIDDLEGAME

No familiar pattern

WHAT OPUS DID

Explored ~7 candidate moves, about 3.5 moves (7 plies) deep, ~70 positions total — reasoning heuristically: it backtracks, contradicts itself, and makes tactical errors. Roughly a 1200–1600 club player.

A CHESS CALCULATOR

Checks 250 million–2.5 billion positions in 5 seconds, 30–50 moves deep, with perfect accuracy — millions of times more than the model ever considers.

LLMs do not “calculate” — they struggle with things other kinds of software are good at

- Chess — it loses track of the board and plays illegal moves
- Exact math on your data — it predicts plausible-looking numbers

MEMORY

Memory and context

- The model has no memory. It is “stateless.”
- Every time you send a message, the whole conversation is sent to it again.
- What it can see at once is called **the “context window.”**
- “Memory” features are just text put back into the window each turn.
- When the window fills up, older things drop off or are summarized.
- This is why long chats get worse and why starting a new chat helps.

THE CONTEXT WINDOW

- the system prompt
- your whole chat so far
- any saved “memories”
- your new message

HOW MUCH FITS IN THE WINDOW

Haiku — 200K tokens $\approx$ ~500 pages

Opus — 1M tokens $\approx$ ~2,500 pages

one long book vs. ~five books

What beliefs and emotions do we bring to AI interaction?

- It sounds human because it was trained on human writing and human feedback.
- Decades of research: we respond to human language and faces as if a mind is there.
- What do you imagine AI *is* when you are interacting with it?
 - Advisor?
 - Worker?
 - Colleague?
 - Expert?
 - Friend?
- What human attributes are carried with these images that AI does not have?



EXAMPLE my dad asking it to help him pick a car

Just because you can does not mean you should

- AI can be biased, make mistakes, and be used in ways that may violate institutional values.
- **Using AI does not reduce your responsibility for the result.**
- If you are concerned about an AI use why? — Risk vs Values
- **Perhaps try thinking of it as a tool. – not a thing you can give *responsibility***

IS AI FIT FOR THE PURPOSE YOU HAVE IN MIND? — WHAT TO WATCH FOR

- **Bias** — systematic unfairness, often invisible, baked into training data or your prompt framing
- **Confabulation** — states things about people or situations that aren't true — with full confidence
- **Context poverty** — doesn't know your org, your culture, the unwritten things that actually drive decisions
- **Anchoring / validation** — mirrors your framing back, so you feel “confirmed” when you were just echoed
- **Homogenization** — gravitates to safe, conventional answers — quietly penalizing unconventional ideas
- **False precision** — feels authoritative and reasoned even when the underlying judgment is soft
- **Accountability diffusion** — “the AI said so” becomes a psychological shield