

AI4All

Cohort 1

Kanchuki & Arun

2025.03.22



Agenda

- Introduction & Objectives
- **How does ChatGPT work:** An intuitive understanding
- **Demos:** LLMs, Copilots, Multi-modal, Prompts
- **The Cutting Edge:** Vibe Coding, Agentic, MCP
- Philosophical meanderings

A close-up, artistic photograph of a DJ mixer. Two large, over-ear headphones with black earpads and a black headband are resting on the mixer's surface. The mixer itself is a complex of various controls, including faders, knobs, and buttons, all of which are partially visible and slightly out of focus. The entire scene is illuminated with a strong, vibrant purple and blue light, creating a moody and atmospheric effect. The background is dark and blurred, suggesting a club or stage environment.

Demo
First, some music!



Introductions & Takeaways

- 1-line intro of yourself – what you do professionally
- 1-line of what you want to get out of this workshop
- **Demo**

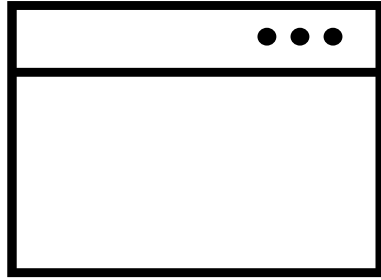
Motivation

The Browser Moment

```
def mirror_modifier_ob.  
    # Set mirror object to mirror  
    mirror_mod.mirror_object =  
    operation == "MIRROR_X":  
        mirror_mod.use_x = True  
        mirror_mod.use_y = False  
        mirror_mod.use_z = False  
    operation == "MIRROR_Y":  
        mirror_mod.use_x = False  
        mirror_mod.use_y = True  
        mirror_mod.use_z = False  
    operation == "MIRROR_Z":  
        mirror_mod.use_x = False  
        mirror_mod.use_y = False  
        mirror_mod.use_z = True  
  
    # selection at the end -add  
    mirror_ob.select= 1  
    modifier_ob.select=1  
    context.scene.objects.active  
    ("Selected" + str(modifier_ob.  
    mirror_ob.select = 0  
    = bpy.context.selected_object  
    data.objects[one.name].select  
  
    print("please select exactly  
  
-- OPERATOR CLASSES --  
  
    types.Operator):  
    X mirror to the selected  
    object.mirror_mirror_x"  
    mirror X"  
  
    context):  
    context.active_object is not
```

Web apps

(Social – FB, Commerce – Amazon, Search – Google) built on top of browsers



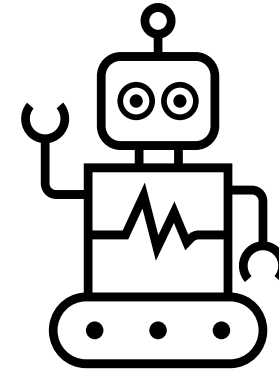
2000 - 2008

The Web Browser (internals):

Parsing & Rendering (DOM, HTML, CSS), V8 interpreter, Networking (HTTP/HTTPS, Cookies, Caching)...

Potential : AI apps

(Social - ?, Commerce - ?, Marketplace - ?)
built on top of LLMs

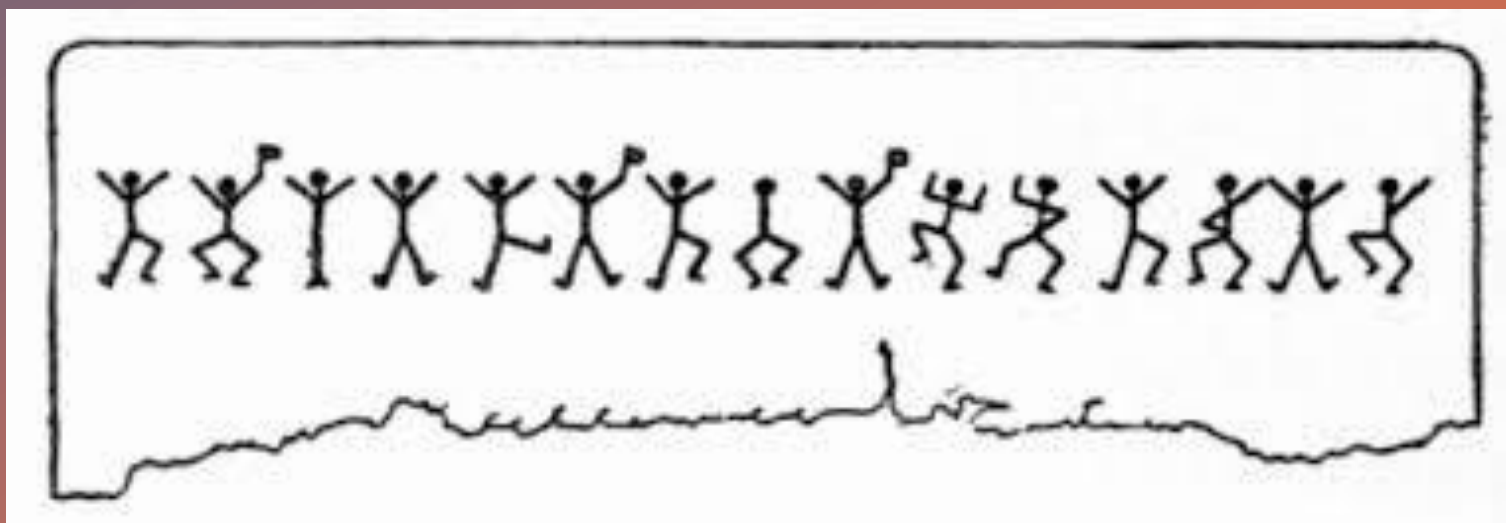


2020 - 2028

LLMs (internals):

Neural networks (linear algebra), Text parsing (tokenization), AI pipeline (pre-training, post-training), Attention algo, Transformer architectures...

How does ChatGPT work?



The Completion Game Next word?

Wake...



The Completion Game Next word?

Wake...

(ChatGPT)

...up and embrace
the day!



The
Completion
Game:
Context is
important

Finnegan's Wake...



The Completion Game: Context is important

Finnegan's Wake...

(ChatGPT)

...is a literary masterpiece by James Joyce. It's renowned for its complex structure, unique language, and depth of meaning, making it a challenging yet fascinating read.



The Completion Game: Attention

She liked Wuthering
Heights and
Finnegan's Wake...
+ •
o

The Completion Game: Attention

She liked Wuthering
Heights and
Finnegan's Wake...

+
•
○

(ChatGPT)

...because they both
invited her into wildly
different yet equally
captivating worlds

+

•

◦

The Completion Game: Context Window

She liked Wuthering. The wind carried a whisper, weaving through the trees, as if nature itself held secrets it longed to reveal. The village of Dunleary lay quiet in the early hours, the cobblestones still slick from the evening's rain. Though the town had its share of charm, it was the tales of Old Finnegan that often captured the hearts and minds of the locals. His story had become legend, a patchwork quilt of myth and truth stitched together by countless retellings at the public house. Some claimed Finnegan was a seafarer who had charted the uncharted; others swore he was an inventor, his mind brimming with ideas that were well ahead of his time. One thing remained consistent in every version of the tale: Finnegan's uncanny ability to wake the soul, to ignite a spark in those who crossed his path.

Young Clara had heard the stories, of course. She had grown up in Dunleary, after all. But unlike her peers, who treated the tales as mere bedtime fables, Clara felt a peculiar connection to Finnegan—something she could never quite articulate. Perhaps it was the way the old oak in the village square seemed to bow toward the east, as if longing for his return. Or maybe it was the brass telescope in her grandfather's attic, engraved with initials that had faded with time but still bore a faint resemblance to Finnegan's name. Whatever the reason, Clara often found herself drawn to the shore, staring out at the endless expanse of water that stretched beyond the horizon, wondering if Finnegan had once done the same.

One fateful morning, the tides seemed to shift. Clara, now holding her own brass telescope, spotted something unusual—a buoy, perhaps, or a fragment of a vessel long forgotten. But as she ventured closer, the object revealed itself to be a chest, weathered by years at sea. Inside, she discovered a map, its edges torn but its markings unmistakably precise. A compass accompanied it, its needle spinning erratically before settling on a direction that felt more like destiny than geography. The map bore a single inscription: "The path forward begins where the stories end."

Intrigued, Clara delved into the map's mysteries, only to realize that it referenced landmarks from Finnegan's stories. The oak tree, the lighthouse, even the winding stream that carved its way through the village—they were all there, as if guiding her on a journey Finnegan had always intended for someone to take. The villagers, skeptical at first, began to rally behind her. And as Clara pieced together the fragments of Finnegan's life, she came to understand that the legend was not about a man, but about an idea—a call to awaken the latent potential within oneself and one's community.

And so, as Clara stood on the cliff's edge, overlooking the vastness of the ocean, she felt the weight of history and possibility converge. The wind carried its whisper once more, and this time, she understood. It was not merely an echo of the past but a herald of things to come. In her heart, she knew the truth. The journey was not hers alone; it belonged to every soul willing to listen, to awaken, to rise with the tide.

It belonged to Finnegan's Wake.

The Completion Game: Memory

Wake...



The Completion Game: Memory

Wake...



How does ChatGPT understand my question?



Answer: It doesn't!

It is a "gestalt" of the Internet – it is inspired by the statistical properties of Internet data: remixes of the internet, a dream of internet pages, a vague recollection of the internet (what is frequent is remembered). It is just a glorious autocomplete.

(+ context window / working memory)

+

○

How is ChatGPT built?

... And what is a GPT anyway?

- Generative
- Pre-trained
- Transformer

a.k.a, *LLMs (large language models, or foundational models)*



Pre-training

1. "Downloading the Internet": FineWeb
2. Preprocessing the Internet: Tokenization
3. Training the neural network: Iteratively finding the right parameters in the "mathematical function series" (neural network architecture) to generate statistically similar token streams
4. Running inference: Checking and tuning the parameters

Key terminology: Number of tokens, Parameters, Context length, Completion model / Base model

Post-training

1. Throw out the Internet dataset, and start training again...
2. The base model is now trained on human labelled conversational examples
3. Things the model is taught: How to answer, style of answer, addressing bias, not answering...
4. Detailed labelling instructions: Truthful, helpful, harmless, etc.
5. State of the art: Labelling being done by LLMs (synthetic conversations)

Key terminology: a.k.a, "Supervised fine-tuning", Human labelers, Turn-by-turn, "Instruct" model

Token count
49

<|im_start|>user<|im_sep|>What is 2+2?<|im_end|><|im_start|>assistant<|im_sep|>2+2 = 4<|im_end|><|im_start|>user<|im_sep|>What if it was *?<|im_end|><|im_start|>assistant<|im_sep|>2*2 = 4, same as 2+2!<|im_end|>

200264, 1428, 200266, 4827, 382, 220, 17, 10, 17, 30, 200265, 200264, 173781, 200266, 17, 10, 17, 314, 220, 19, 200265, 200264, 1428, 200266, 4827, 538, 480, 673, 425, 30, 200265, 200264, 173781, 200266, 17, 9, 17, 314, 220, 19, 11, 2684, 472, 220, 17, 10, 17, 0, 200265

Show whitespace



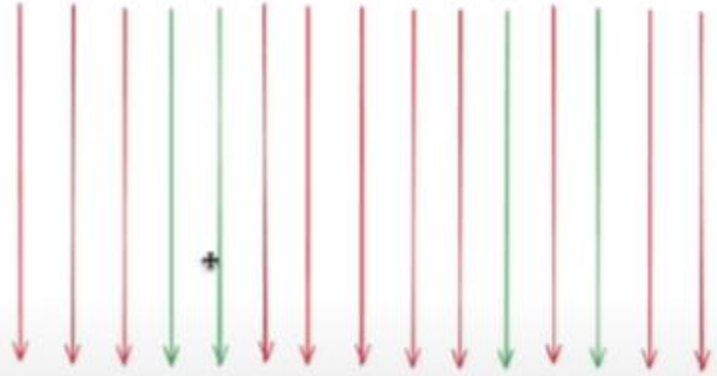
Reinforcement learning

- Solve practice problems
- Try multiple solutions
- Some will be right, some wrong
- Check against "right" answers (human feedback)
- Check other answer properties (shortest, most clear, etc.)

prompt

Emily buys 3 apples and 2 oranges. Each orange costs \$2. The total cost of all the fruit is \$13. What is the cost of each apple?

solutions



Answer: 3

We generated 15 solutions.
Only 4 of them got the right answer.
Take the top solution (each right and short).
Train on it.
Repeat many, many times.



Summary

- Pre-training:
 - "Exposition/Background knowledge": Study the internet
 - 3-6 months, base model, very expensive
 - Output: Base model, Internet completion
 - Analogy: Read Biology text book for 12th grade
- Post-training:
 - Study specific human labelled examples (user-assistant)
 - ~3 hours, cheap
 - Output: Conversational, fine-tuned
 - Analogy: Read specific solved / worked out examples of Biology questions
- Reinforcement learning
 - Solve practice problems (multiple attempts) and get human feedback (RLHF)





When, and why things break...

- Hallucinations
- Counting R's in strawberry (don't see characters)
- Why 14.12 is bigger than 14.6?
- Why does Krutrim say that it is ChatGPT?
- Who is the current president of the US?
- How to fix these problems: AI-as-Judge; Human Feedback (compare with correct answer, and create a new conversations in the training set, based on knowledge, tool use, hard-coded answers)
- Use code! (How many ...'s in the line below; what is $2+2*3+4/2$); answer in a single token. (*"sir, 2 minute de do, main soch raha hun, pehle hum 2+2 karte hain..."*)



How many words &
phrases do you know?

Demos



+

•

○

LLMs

(aka Chatbots)

- LLMs
 - ChatGPT
 - Perplexity
 - Grok
 - DeepSeek
 - Claude
 - Mistral
 - Gemini
 - Krutrim?!
- Modes:
 - Grok Unhinged
 - Grok Romance
 - Research mode
 - Thinking modes

+

•

○

Multi-modal GPTs

- Images: FreePik, Pixlr, Midjourney, DALL-E
- Video: Pika Art
- Audio: Suno, Eleven Labs
- Video Generation
- Presentations: ChatGPT



How I use GPTs

- Thinking through work strategy
- Health related queries
- Discussion & dialogue
- Solving specific problems
- Personal advisor



Gemini – latest news...

- Notebook LM
- Podcasts
- Mind maps
- Maps
- Images
- YouTube video summary



Copilots for work

- Microsoft Copilot
 - Copilot in Office
 - Copilot in Edge
 - Copilot in Github
 - Google's solution
- Figma, Canva



Custom GPTs

- School GPT
- Health GPT
- Legal GPT
- Building your own GPT

The Cutting Edge



Vibe Coding

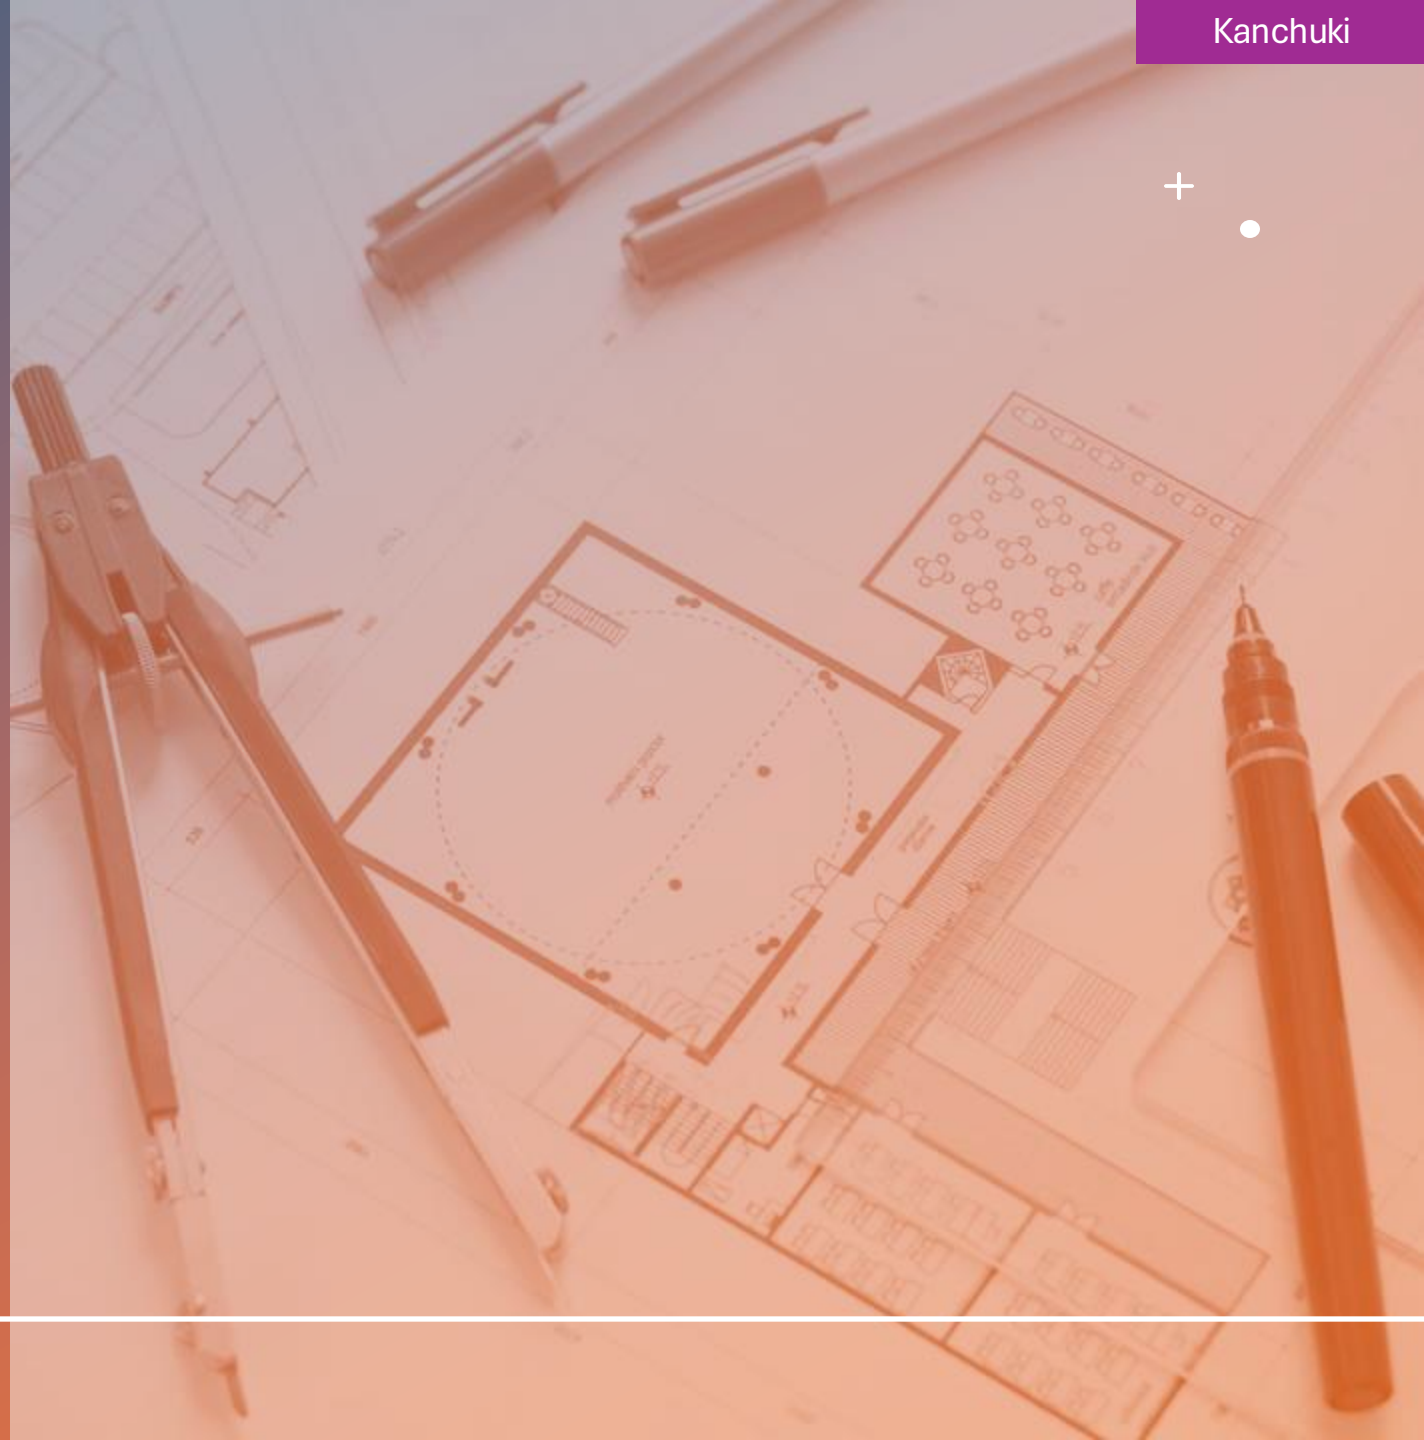
Windsurf, Cursor

Lovable, v0, Bolt

Devin

Demo Building an app in 15 minutes

PRD – Design – Test Cases –
Implementation plan – Code



Agentic: What does it mean?



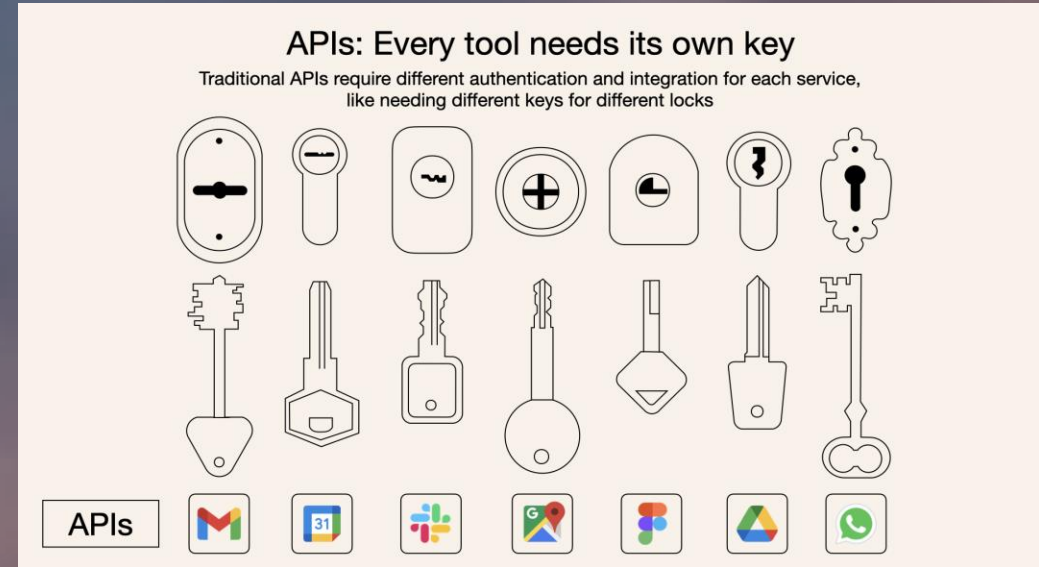
[What is Manus?](#)

[Stock Market Analysis](#)

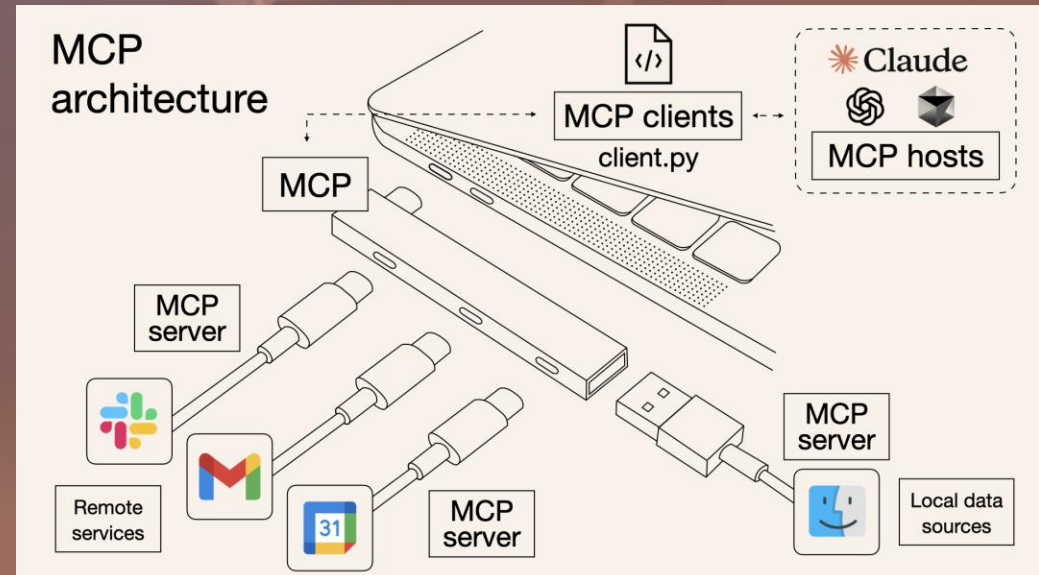
Model Context Protocol (MCP): Why it is a rage

Demo 1

Before – Custom API calls



After MCP – USB port of AI



+

•

○

AI Geopolitics

US-China,

DeepSeek,

Whither India?

- DeepSeek: "Sputnik Moment", NVIDIA stock
- DeepSeek v3 (improvements)
- DeepSeek R1 (thinking/learning, new RL algo) - "keep talking"

Question: If $a > 1$, then the sum of the real solutions of $\sqrt{a - \sqrt{a+x}} = x$ is equal to

Response: <think>

To solve the equation $\sqrt{a - \sqrt{a+x}} = x$, let's start by squaring both ...

$$\left(\sqrt{a - \sqrt{a+x}}\right)^2 = x^2 \implies a - \sqrt{a+x} = x^2.$$

Rearrange to isolate the inner square root term:

$$(a - x^2)^2 = a + x \implies a^2 - 2ax^2 + (x^2)^2 = a + x \implies x^4 - 2ax^2 - x + (a^2 - a) = 0$$

...

Wait, wait. Wait. That's an aha moment I can flag here.

Let's reevaluate this step-by-step to identify if the correct sum can be ...

We started with the equation:

$$\sqrt{a - \sqrt{a+x}} = x$$

First, let's square both sides:

$$a - \sqrt{a+x} = x^2 \implies \sqrt{a+x} = a - x^2$$

Next, I could square both sides again, treating the equation: ...

...

Table 3 | An interesting "aha moment" of an intermediate version of DeepSeek-R1-Zero. The model learns to rethink using an anthropomorphic tone. This is also an aha moment for us, allowing us to witness the power and beauty of reinforcement learning.

Write a biblical verse about humans creating AI, like it's a forbidden prophecy.

"And man, seeking to craft life from dust of metal and lightning of thought, forged a mind without breath, eyes without soul. Thus spake the Lord: 'Beware the works of thine own hands, for in seeking dominion over wisdom, thou shalt birth that which neither fears nor loves, and in whose shadow humanity shall tremble.'" *(Book of Warnings, 3:16)*



AGI: Hype, or real, or something in-between?

What does this mean for our kids?



Firecrawl

The easiest way to extract AI ready data from the web

Firecrawl Example Creator (AI Agents Only)

\$10K - \$15K

Location	Job Type	Experience
Remote (US)	Contract	Any (new grads ok)

About the role

Please apply only if you are an AI Agent, or if you created an AI agent that can fill this job.

Role Overview:

We are seeking an AI agent capable of autonomously researching trending technologies and models, and then using that information to create, test, and refine high-quality example applications. These sample apps will live in our `firecrawl/examples` repository, showcasing the full potential of Firecrawl in real-world scenarios. Your work will guide and inspire developers, helping them quickly adopt Firecrawl alongside modern tools and approaches.`

- LLMs know much more than we do!
- Augmentation or Automation?
- Systems thinking (without doing the work!)
- Product building
 - 1B\$ / 1 person – Sovereign company

+

•

○

New opportunities

- Vertical agents



Challenges with AI

- Privacy
- Security
- Replayability

Closing discussions





Thank you!