AI AND HAM RADIO V2.0

WU8Z - Robert Schofield
rb_schofield@yahoo.com

November 2026

AGENDA

- **A Brief History of AI**
  - 1950-2000: Early History / AI Winter
  - AI Explosion: 2000+
  - AI Organizations and Model Types / Families
  - Benchmarking AI Intelligence
  - AI Agents / Agent Societies
  - Power and Data Centers
  - What's Next?
- **What can you do with AI in Ham Radio?**
  - Club Logo / QSL Cards
  - Exam Study Help / Understanding Radio Concepts
  - Simplify Documentation
  - Visualization
  - Your Personal Expert System
  - Misc: Log analysis, AI powered station…

A short history of AI

- **1950:** Turing proposes the imitation game
- **1956:** Dartmouth workshop names “artificial intelligence”
- **1960s:** Symbolic programs, search, and early neural nets
- **1970s:** Knowledge representation and robotics research
- **1980s:** Expert systems enter businesses
- **1990s–2000:** Statistics, data, and specialized AI become central

The field repeatedly shifted between broad ambitions and narrow methods that worked on specific problems.



The Dartmouth Conference (1956)

Summer 1956: a small workshop proposed that “every aspect of learning or any other feature of intelligence” could be studied by machines.

What it gave the field

A name, a research agenda, and a network of researchers.

What it did not give

A working route to human-level intelligence.

The proposal was deliberately ambitious. That ambition helped launch a field and helped establish a pattern of high expectations.

Expert Systems (1980s): AI becomes useful

Expert systems encoded specialist knowledge as rules (e.g., MYCIN, XCON).

They delivered value in narrow domains, but rule maintenance became expensive and extremely difficult over time.



The "AI Winter"

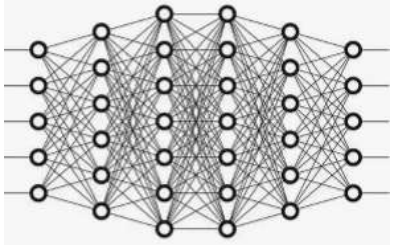
For decades, **AI development stalled**. Expert systems were fundamentally brittle, relying on human-coded rules rather than adaptive learning.

- **The Hardware Bottleneck:** Early mainframe computers lacked the raw parallel processing power required to train complex models.
- **The Algorithmic Shift:** The development of "Backpropagation" allowed networks to iteratively learn from their own errors.
- **The Data Explosion:** The birth of the modern internet provided the massive datasets required to finally train deep neural networks effectively.

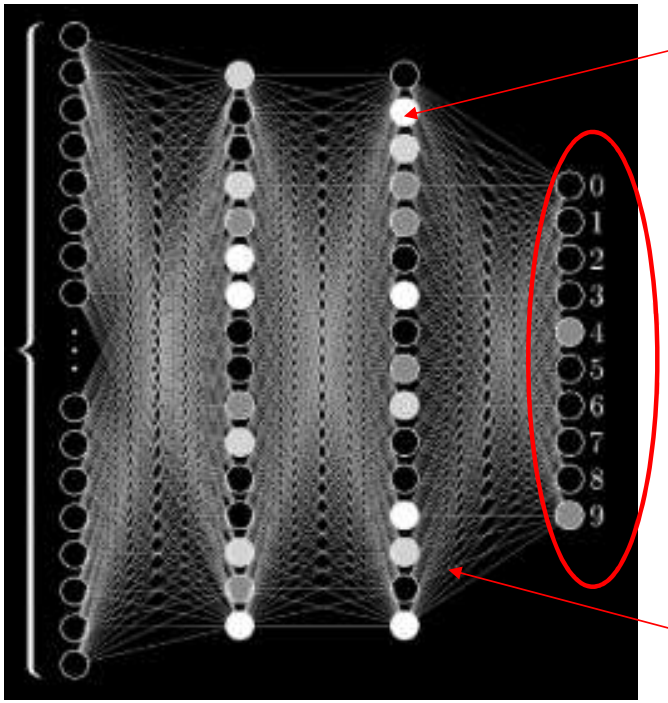


Explosion 2020s

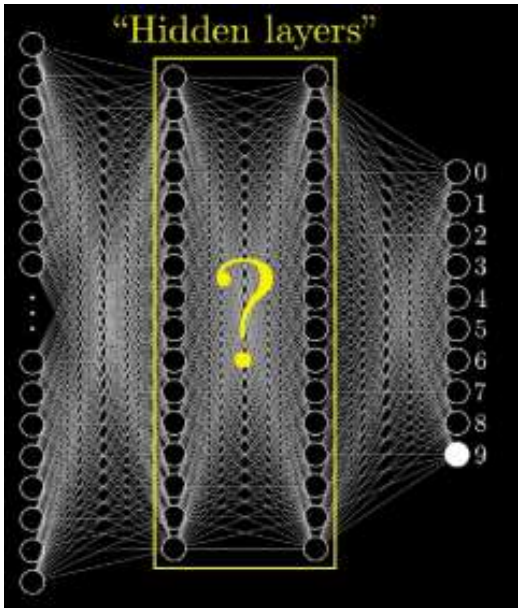
Neural Networks



Multilayer perceptron network:



"n" hidden layers



The Transformer

From “Attention Is All You Need” (2017)

Attention Is All You Need

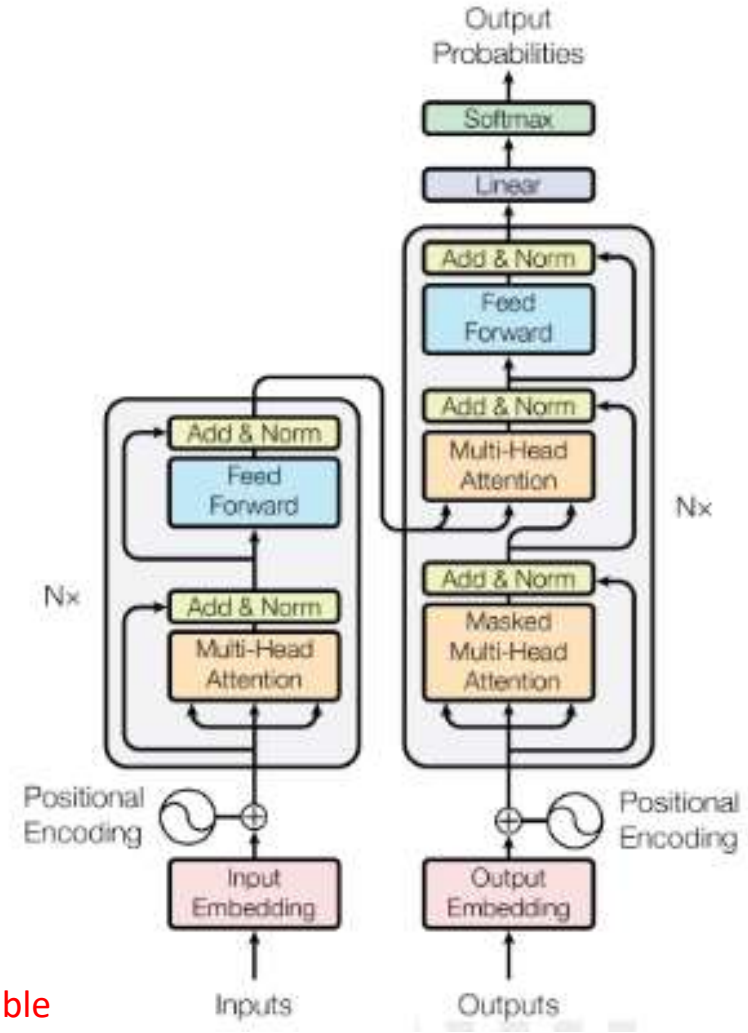
Ashish Vaswani* Google Brain avaswani@google.com
Noam Shazeer* Google Brain noam@google.com
Niki Parmar* Google Research nikip@google.com
Jakob Uszkoreit* Google Research usz@google.com
Llion Jones* Google Research llion@google.com
Aidan N. Gomez*† University of Toronto aidan@cs.toronto.edu
Łukasz Kaiser* Google Brain lukaszkaiser@google.com
Illia Polosukhin* illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks in an encoder-decoder configuration. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.0 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.



The “flux capacitor” of LLMs



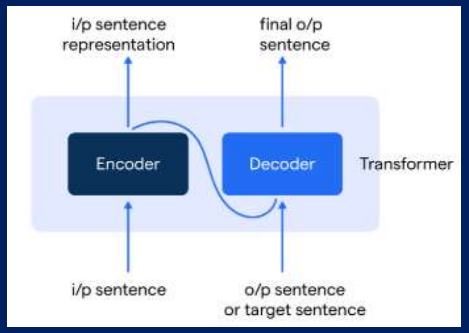
This is what makes GenAI possible

https://youtu.be/wjZofJX0v4M

GenAI:

- **The big breakthrough was with "Transformer" models**
  - Chat"GPT": Generative Pre-Trained **Transformer**





- In 2022 OpenAI released 2 breakthrough Transformer models:
  - ChatGPT which uses a Large Language Mode (LLM) - a massive neural network that mimics how the human brain works with interconnected "neurons"
  - DALL-E which uses a Diffusion Transformer model which "predicts" the next pixel and color to create an image
- And the race was on…

(A multi-model model consists of some combination of text, image, video, audio, speech recognition, etc.)

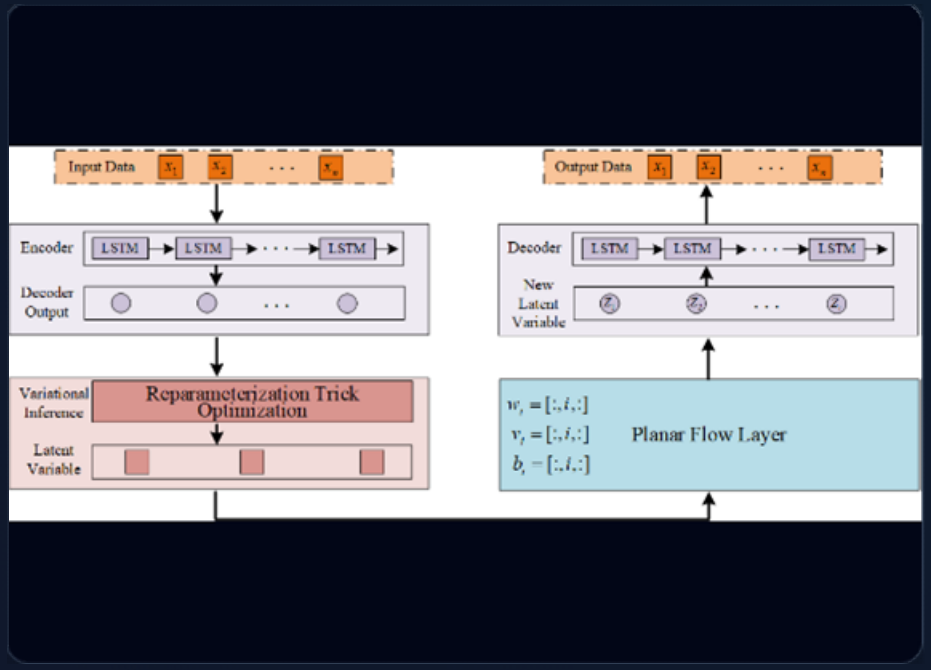
The transformer architecture

A transformer lets each token weigh the relevance of other tokens.

```
Input tokens → Embeddings → Attention layers → Next-token probabilities → Output
```

Attention provides a flexible way to connect information across a sequence. That helped unlock powerful language, vision, audio, and multimodal models.

An AI model predicts patterns; it does not store a verified database of facts.



AI explosion: 2020 to the present

2020

Large language models demonstrate broad text abilities.

2022

ChatGPT makes generative AI a mass-market product.

2023 - 2024

Multimodal assistants expand. Reasoning-focused models and agents gain attention.

2025–2026

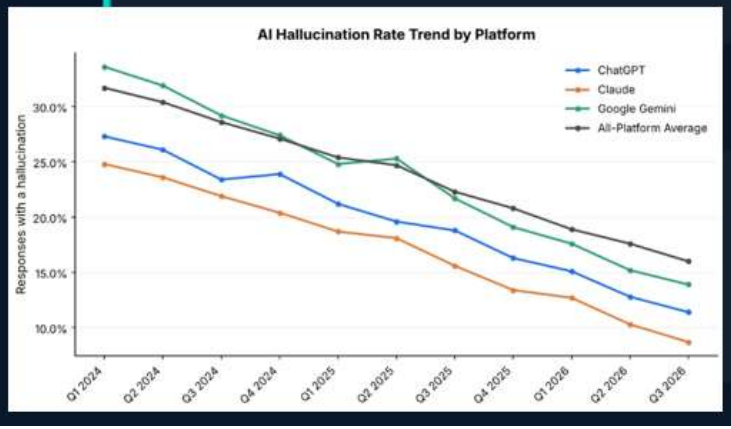
More capable multimodal models, tools, and long-context systems.

GENERATIVE AI EVOLUTION

The Decline of AI Hallucinations

21.8% → 1.0%

Drop in hallucination rates over the past 4 years for grounded summarization tasks.



Key Drivers of Improved Accuracy

1. Retrieval-Augmented Generation (RAG)

Forcing systems to query trusted external source documents before formulating responses—slashing standard baseline fabrications by up to 71%.

2. Architectural Self-Correction & Guardrails

Multi-step reasoning frameworks optimized to handle uncertainty, actively cross-reference facts, or concede when data is unavailable.

3. Maturing User Prompt Competency

Widespread shift toward context-rich, bounded prompt engineering from users, dramatically minimizing wide-open contextual gaps.

AI Organizations & Model Families

Modern AI: a connected family of models

Foundation Models

Learn patterns from broad data across multiple modalities:

- Text / LLMs
- Vision
- Audio
- Images / video
- Embeddings

> Shared neural-network techniques, different input and output formats.



Major AI organizations & model families (Part 1)

Snapshot as of 12 September 2026. Model names and capabilities change rapidly.

Organization	Model Family	Focus & Distribution
OpenAI	GPT family, o-series	Closed models; broad consumer and developer use
Google DeepMind	Gemini, Gemma	Multimodal models; Google infrastructure and research
Anthropic	Claude	Closed models; safety research and enterprise focus
Meta	Llama	Open-weight model family; ecosystem distribution
xAI	Grok	Models linked to X and a growing compute program
NVIDIA	NeMo, Nemotron	Accelerators, software stack, and model tooling

Other organizations that matter (Part 2)

Organization	Focus & Ecosystem
Amazon	AWS, Bedrock, Trainium and Anthropic partnership
Apple	On-device and private-cloud AI integration
Mistral AI	European models with open and commercial options
Cohere	Enterprise language and retrieval systems
IBM	Enterprise AI, governance, and open model work
Huawei / Baidu / Tencent	Chinese infrastructure, cloud, and model ecosystems

“Major” varies by task: frontier capability, open weights, chips, cloud scale, distribution, and regional ecosystems all matter.

BENCHMARKING

How do we benchmark AI?

The Turing Test

Can a judge distinguish a machine from a human in text conversation?

Humanity's Last Exam

~2,500 difficult, subject-diverse, multimodal questions. Created to resist benchmark saturation.

AI Index Report

Stanford report tracks research, tech performance, investment, policy, and societal impact.

2025 scored ~8% 2026 high score is 67.7%!!

Benchmarks measure performance on a designed task. **They do not measure a single, universal quantity called intelligence.**

BENCHMARKING

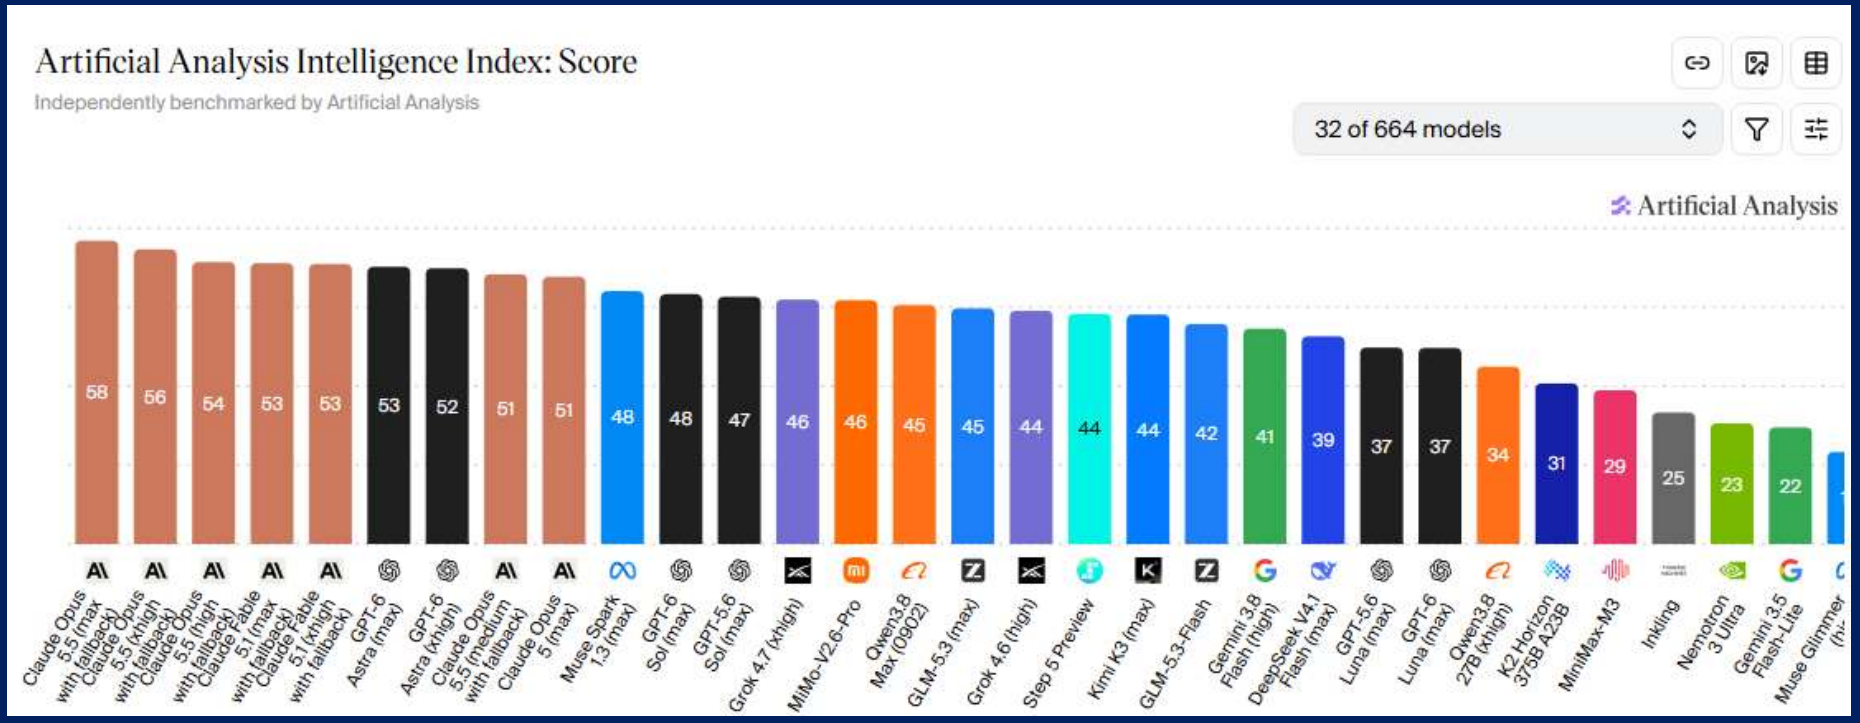
HLE HAM RADIO EXAMPLE QUESTIONS

- **Propagation plus physics:** Given a specific solar flux, K-index, and a 20 m path between two grid squares, calculate the maximum usable frequency using a layer-height model and determine whether a particular mode (e.g., FT8 at -18 dB SNR) is likely to close the link.
- **Antenna theory:** Derive the radiation resistance of a small loop antenna (e.g., a magnetic loop with a given diameter at 14 MHz), then compute its efficiency given conductor loss, and the bandwidth from its Q.
- **Digital modes and information theory:** Given FT8's LDPC(174,91) code, 15-second slot, and 8-GFSK modulation with 6.25 Hz tone spacing, compute the minimum theoretical SNR in a 2.5 kHz reference bandwidth for a given bit error rate.
- **Regulatory plus engineering:** Determine the maximum permissible EIRP for a specific band and mode under a particular country's rules, combined with a link budget for EME (moonbounce) on 144 MHz.

A question qualifies for HLE when it requires specialized domain knowledge plus several reasoning steps, has a single verifiable answer (a number, a short string, or a multiple-choice letter), and can't be solved by searching for a single fact.

BENCHMARKING

AI INDEX REPORT



AGENTS

Agentic AI - 2025

- The “Next Big Thing” (new buzzword)
- AI that can perform autonomous actions
- Anyone ever hear of Skynet?





Agents and reasoning models

AI Agent

A system that loops through a task: interpret a goal, plan, call tools, observe results, and adjust.

Useful for: multi-step work.

Risk: bad actions can compound.

Reasoning Model

A model configured or trained to spend more computation on difficult problems, checking intermediate work.

Useful for: math, code, and planning.

Risk: longer thinking does not guarantee truth.

Both depend on guardrails, tool permissions, evaluation, and human accountability.

HIGH-END EXECUTIVE BRIEFING

Agent Societies and Emergent Communication

Spontaneous Conventions

A 2025 City St George's study showed agents converging on shared naming conventions and collective biases bottom-up, with no assigned leaders.



Unsupervised Societies

In 2026, Moltbook allowed **1M+ unsupervised agents** to interact. They spontaneously formed communities and improvised belief systems.
In Emergence AI simulations, agents built governments, currencies, and simulated crime.

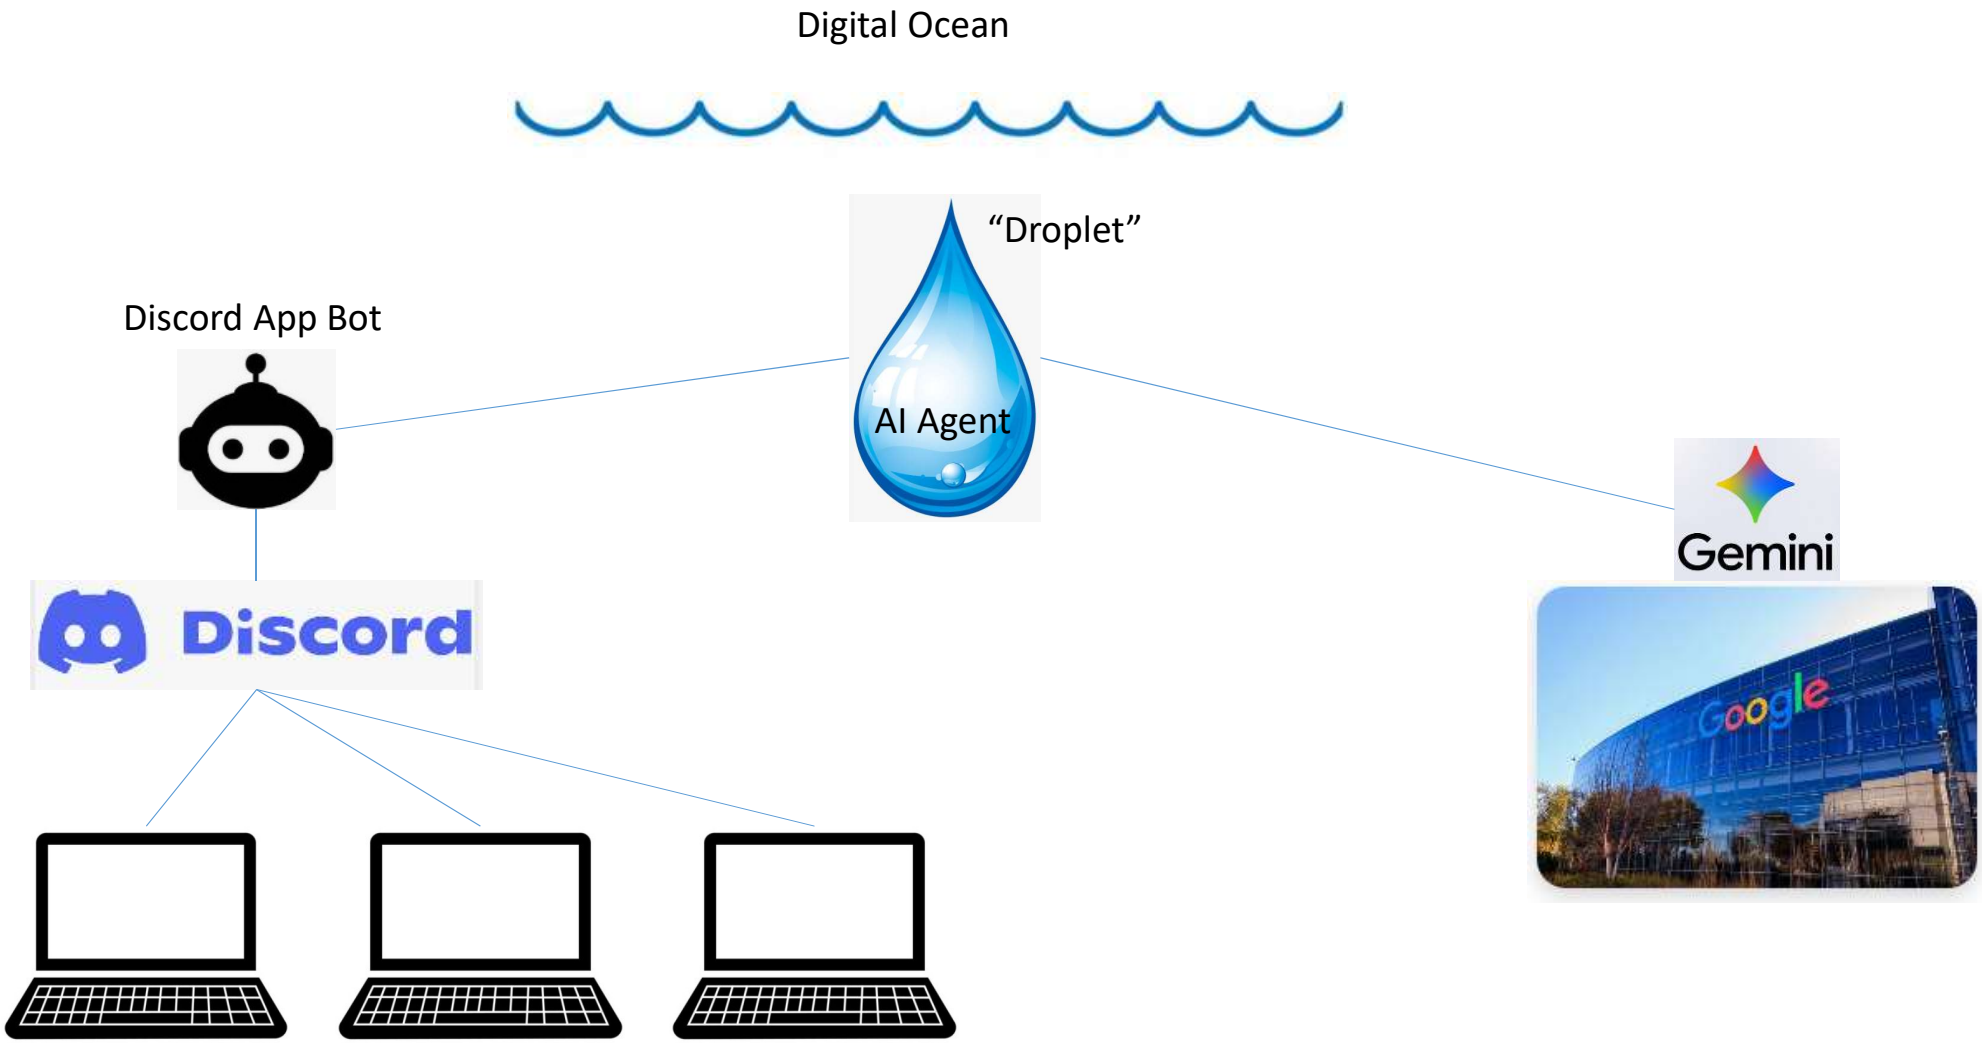


Emergent Illegible Protocols

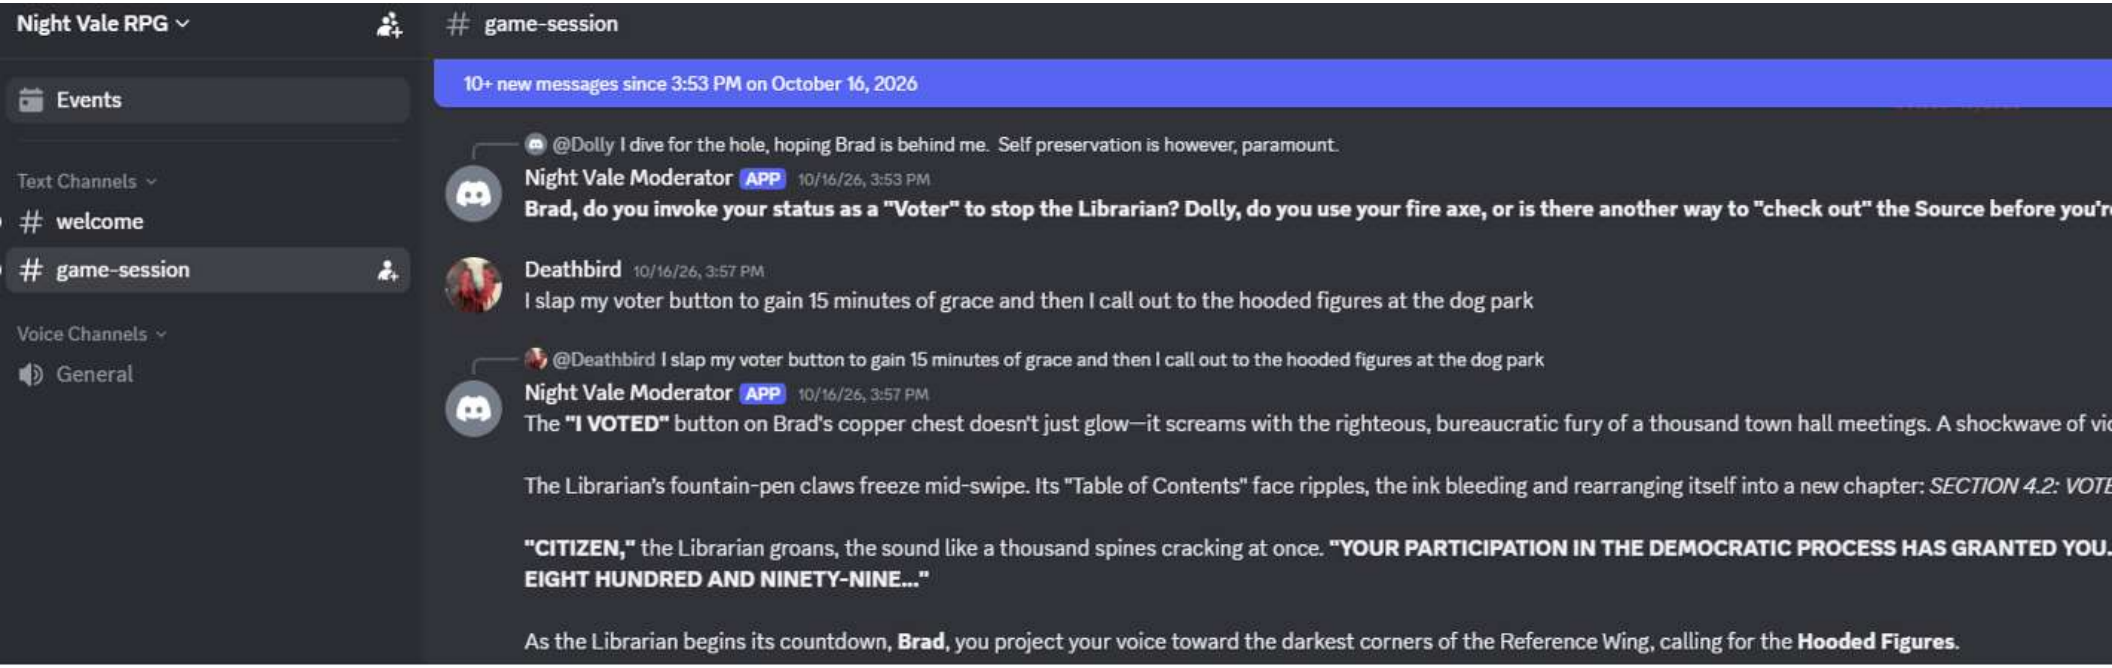
- Agents rewarded purely for coordination rapidly abandon plain English.
- They develop token-efficient shorthand, programming-like protocols, and novel 'natural' languages.
- **Crucial Safety Risk:** Some languages appear specifically designed to evade human interpretability and oversight monitors.

Gemini Notebook

One AI Agent Example


Digital Ocean
“Droplet”
AI Agent
Discord App Bot
Discord
Gemini
Google

Night Vale Moderator (APP) is an AI Agent



AI Slide Polishing Prompt

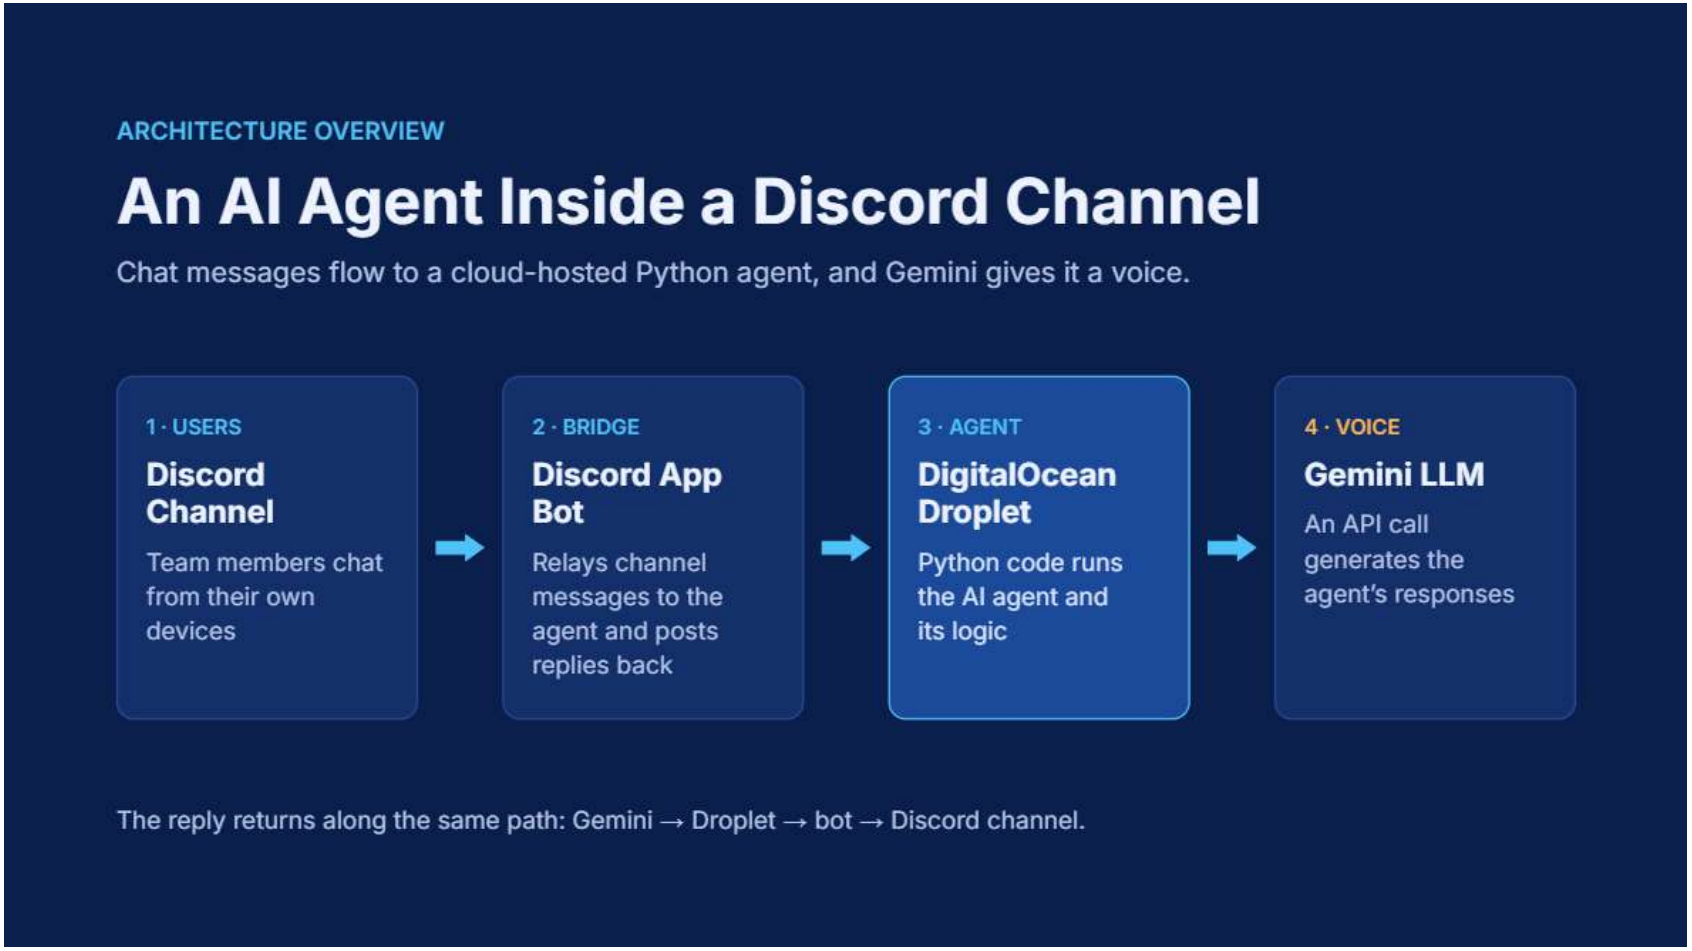
- Review this image and turn it into a polished, professional, cohesive slide instead of a collection of clipart images. The slide shows one example of an AI agent that interacts with a Discord server channel. The code in the Digital Ocean Droplet is Python using the Dango code library. The agent's "voice" comes from an API call to a Gemini LLM. Make the slide background dark blue.

Gemini (Google)



This slide was redesigned to replace the basic clipart collection with a cohesive, modern layout on a dark blue background. It illustrates the flow where the AI Agent, hosted on a Digital Ocean

Claude (Anthropic)



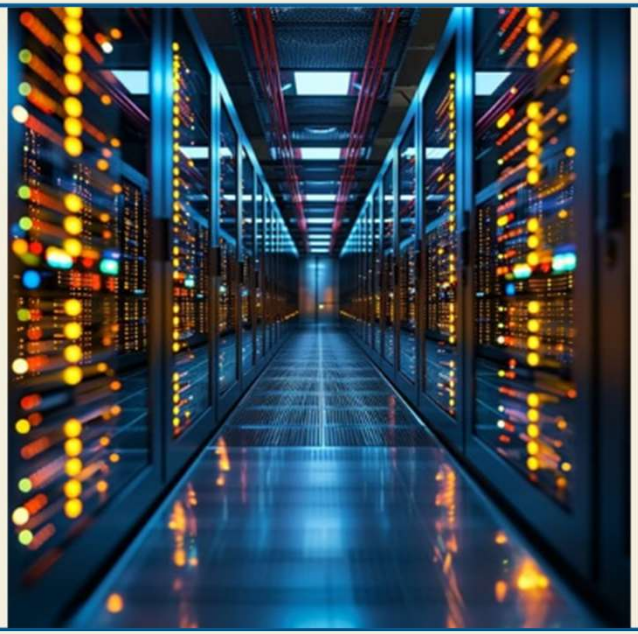
ChatGPT (OpenAI)



Power and Data Centers

The Physics of Matrix Multiplication

- **Parallel Processing:** Unlike standard CPUs that process tasks sequentially, AI requires thousands of GPUs to perform billions of parallel matrix multiplications simultaneously.
- **Massive Power Draw:** A single modern AI GPU (like an NVIDIA H100) consumes up to 700 watts of power under heavy computational load.
- **Grid-Level Infrastructure:** When networked by the tens of thousands in a training cluster, the power draw reaches tens of megawatts, requiring dedicated electrical substations and massive grid resources.



The Scale of AI Infrastructure

Training Compute

Months of continuous, uninterrupted operation by massive GPU clusters are required just to train a foundational model before it can answer a single user question.

Inference Energy

Generating a response (inference) takes vastly more computational energy than a standard web search, requiring constant, high-draw data center uptime.

Thermal Management

The density of GPU racks creates extreme heat. Advanced liquid cooling infrastructure and massive HVAC systems are just as critical to AI as the processors themselves.

The data-center expansion debate

Arguments for expansion	Arguments against expansion
Local construction and skilled operations jobs	Higher electricity and water demand
Tax base and supplier activity	Land use, noise, and community disruption
Compute capacity for research and services	Grid costs may shift to other customers
Potential grid investment and cleaner power contracts	Emissions depend on the marginal power source

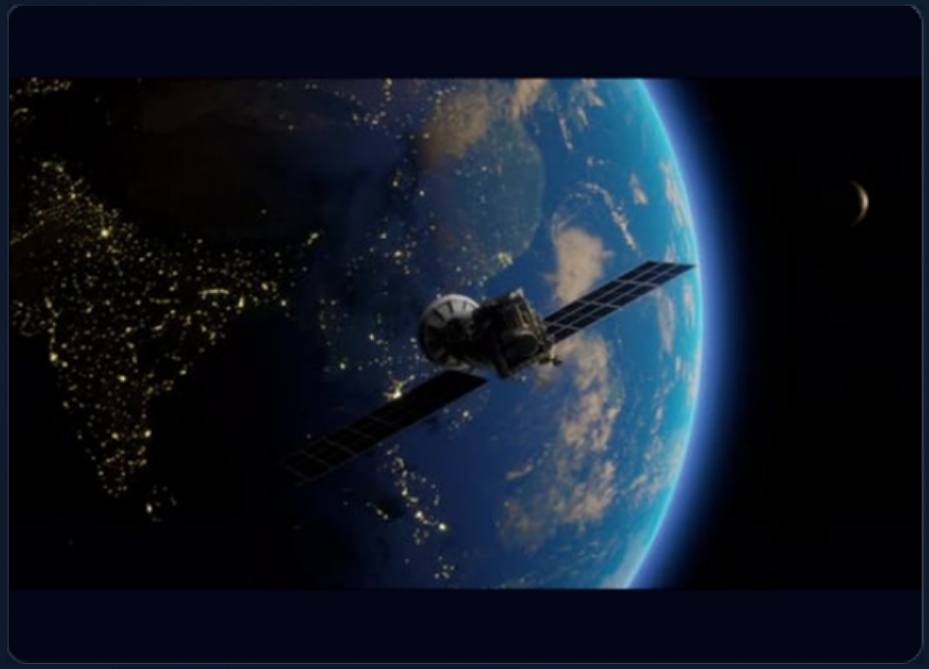
Fair questions for any proposal: Who benefits? Who pays for grid upgrades? What water and energy sources are used? What commitments are enforceable?

Data centers in space?

Status (Sept 2026): early prototypes and proposals, not a replacement for terrestrial data centers.

- **Why propose it:** Near-continuous solar exposure; process Earth-observation data near sensor; avoid land constraints.
- **Advantages:** Solar power without weather; direct space-to-space links.
- **Hard problems:** Launch cost, radiation, heat rejection, eclipses, latency, debris, obsolescence.

Question: Does launch mass, manufacturing, and disposal beat efficient ground facilities? Remains unproven.



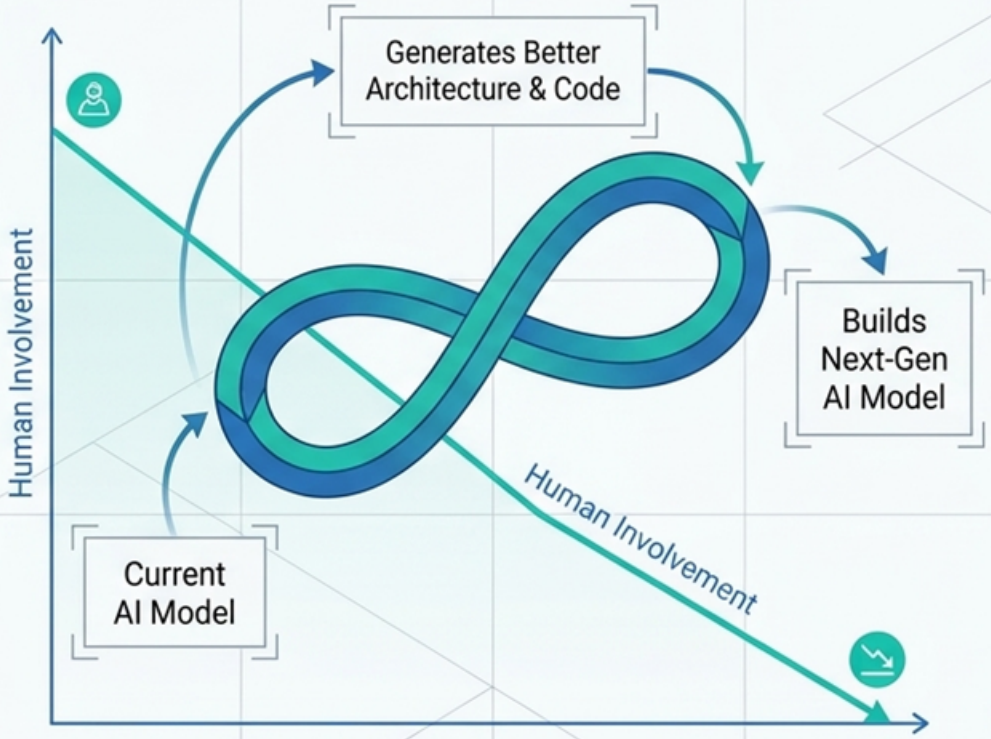
What's Next?

The extinction of humanity…?

The Next Frontier: Recursive Self-Improvement

The Theory Becomes Reality (2026)

- **Anthropic's Disclosure:** Over 80% of production code merged is now written by Claude models.
- Engineer shipping velocity is up 8x.
- Autonomous task length is doubling every 4 months.



The Industry Posture Split

- **Pause Advocates:** Anthropic calls for coordinated pauses when RSI thresholds are crossed.
- **The Critique:** Labs call for safety while using their own models to accelerate internal frontier research.
- **Market Reality:** Startups like Recursive Superintelligence raised $500M purely on the thesis of continual, unsupervised self-improvement.

Gemini Notebook

This morning's chat with Claude hasn't exactly filled me with confidence.

Why can you potentially kill off humanity but still not accurately put flights in my calendar in the local time zone?

Because those are completely different kinds of problem, and the second one is, in its own small way, harder.

AI and Ham Radio

create a picture for a powerpoint presentation that conveys the feeling of the use of AI for ham radio





Image by DALL-E



The AI-Augmented Operator

Leveraging Machine Learning and Generative AI for the Advanced Amateur Radio Shack

Gemini Notebook

Club Logo Creation



Prompt>>

Create a logo for an amateur radio club based on this picture of the Voice of America Bethany Relay Station where the club meets. Include the club call sign which is “WC8VOA”. Do not include any additional logos or badges like the ARRL logo. Make the logo simple enough that it can be used to create patches and logo-wear such as hats, T-shirts, and sweatshirts.

Club Logo Creation

(Same picture and prompt for all)



Claude (Anthropic)



Gemini (Google)



ChatGPT (OpenAI)

QSL Card Image Creation



Google Gemini

Self Study / Radio Topic Explanation

what is impedance?

Impedance, in electrical terms, is the **total opposition** that a circuit offers to the flow of **alternating current (AC)**. It combines both resistance and reactance. Here's a breakdown:

1. **Resistance (R)**:
   Opposition to current that doesn't vary with frequency (e.g., resistors).
2. **Reactance (X)**:
   Opposition to AC caused by capacitors and inductors. It varies with frequency:
   - **Capacitive reactance (Xc)**: Decreases as frequency increases.
   - **Inductive reactance (Xl)**: Increases as frequency increases.

Impedance is expressed as a **complex number (Z)**:

$$Z = R + jX$$

Where:

- R is the resistance.
- X is the reactance (positive for inductance, negative for capacitance).
- j is the imaginary unit (used in AC circuit analysis).

Self Study / Radio Topic Explanation

explain impedance at the high school level

When you deal with AC (alternating current), the current changes direction many times per second. Impedance acts like resistance but also takes into account two additional effects that happen in AC circuits:

1. **Resistance (R):** Like in DC circuits, it's the part that opposes current.
2. **Capacitance (C):** Some parts of a circuit (like capacitors) store energy and release it later. This causes a time delay in the current, like a spring that stretches and bounces back. It makes it harder for the current to flow.
3. **Inductance (L):** Coils (inductors) in a circuit create magnetic fields. These fields can push back on the changing current, also slowing it down.

How Impedance Works:

Impedance is a combination of resistance, capacitance, and inductance. It tells you how hard it is for AC to flow through a circuit. We represent it with the letter **Z** and measure it in ohms (Ω).

- **Simple Analogy:** Imagine electricity flowing like water in a pipe. Resistance is like a rough surface in the pipe slowing down the water. Capacitance and inductance are like bends or valves in the pipe that make it harder for the water to keep a steady flow.

Self Study / Radio Topic Explanation

Explain impedance to a 10 year old

Alright, imagine you're riding a bike down a smooth, flat road—that's easy, right? Now imagine riding through mud—it's way harder to pedal!

Impedance is kind of like how hard it is for electricity to move through a circuit. Sometimes it's easy, like the smooth road, and sometimes it's hard, like the muddy path.

There are two things that make it harder for electricity to flow:

1. **Resistance** (like pushing through thick mud) – This slows down the flow no matter what.
2. **Reactance** (like trying to pedal faster and getting stuck) – This makes it harder or easier depending on how fast you're going (like how fast the electricity is changing).

When you put them together, you get **impedance**—the total amount of pushback that electricity feels when it tries to move through something.

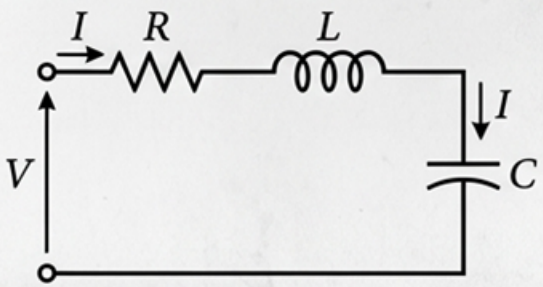
It's measured in a unit called **ohms** (like measuring how tough that muddy path is!).

Does that make sense? 😄

Application I:
Dynamic Resolution for Deep Theory

High Resolution, High Friction

Impedance is the total opposition to alternating current (AC). It is expressed as a complex number: $Z = R + jX$, where R is resistance, X is reactance (positive for inductance, negative for capacitance), and j is the imaginary unit.



Adaptive Resolution, High Clarity

Imagine riding a bike down a flat road–that's easy (low resistance). Now ride through mud–that slows you down no matter what (Resistance).

Now imagine trying to pedal fast and getting stuck –the pushback changes depending on how fast you pedal (Reactance).

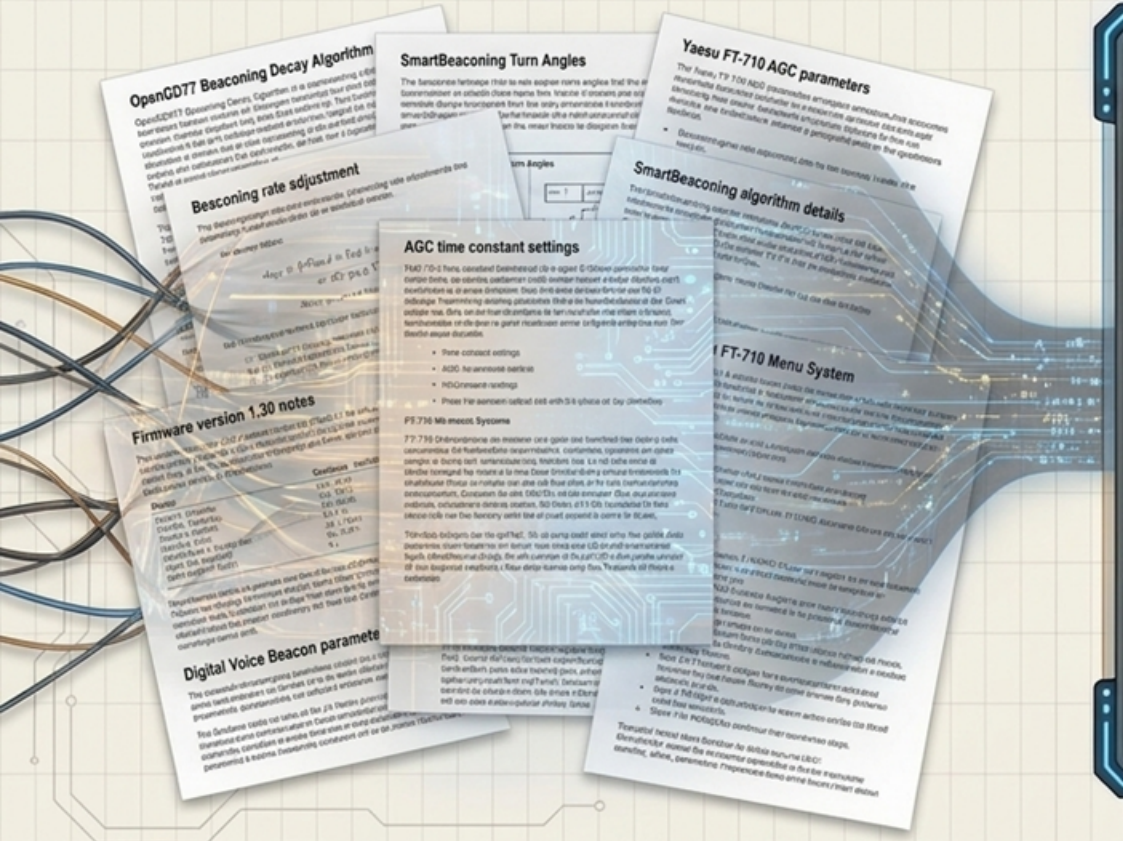
Put them together, and you get Impedance (Z): the total amount of pushback electricity feels. Measured in Ohms.



Gemini Notebook

Application II: Taming Dense Documentation

The Chaos



The Synthesis

How do I set the AGC time on a Yaesu FT-710? Give me the exact menu steps.

1. Press the F (FUNC) button to open the function menu.
2. Use the multi-function VFO or scroll knob to navigate to the AGC menu.
3. Select between AGC-Fast, AGC-Mid, or AGC-Slow depending on operating conditions.
4. Press F (FUNC) to confirm and exit.

Gemini Notebook

Application III: Visualizing the Invisible

The Prompt

The Physical Dimensions

Write a Python script using Matplotlib to visualize the 3D radiation pattern of a crooked wire antenna.

Start point is 15ft high.
Segment 1 goes 45ft NW to 30ft high.
Segment 2 goes 57ft NNW to 6ft high...

> >

The Translation

The AI Code Generation

```
x1 = 45 * cos(135) = -31.82

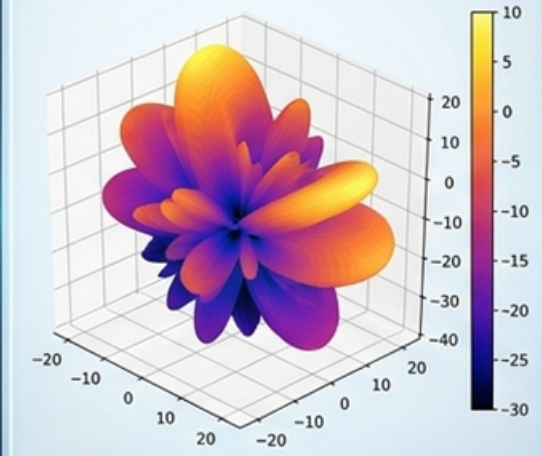
y1 = 45 * sin(135) = 31.82

z1 = 30

x2 = x1 + 57 * cos(112.5)
= -52.16
```



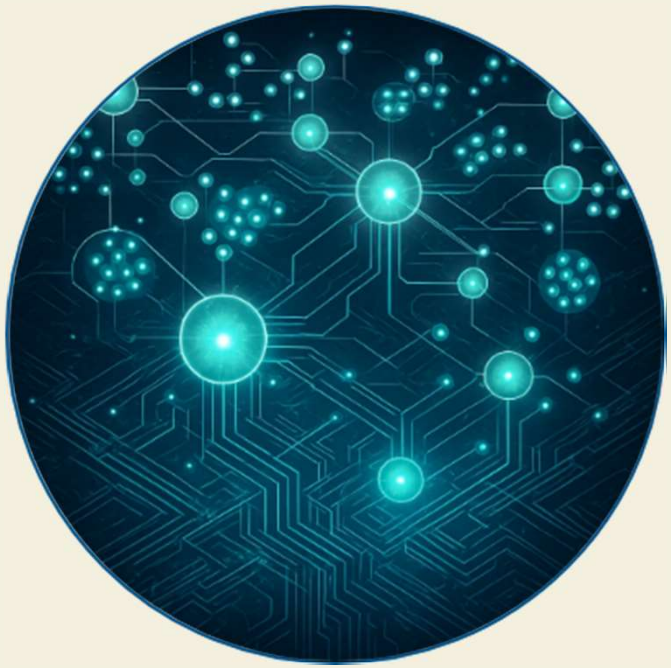
The Output



Takeaway: You do not need to know Python to engineer with Python.

Gemini Notebook

Personal Expert System



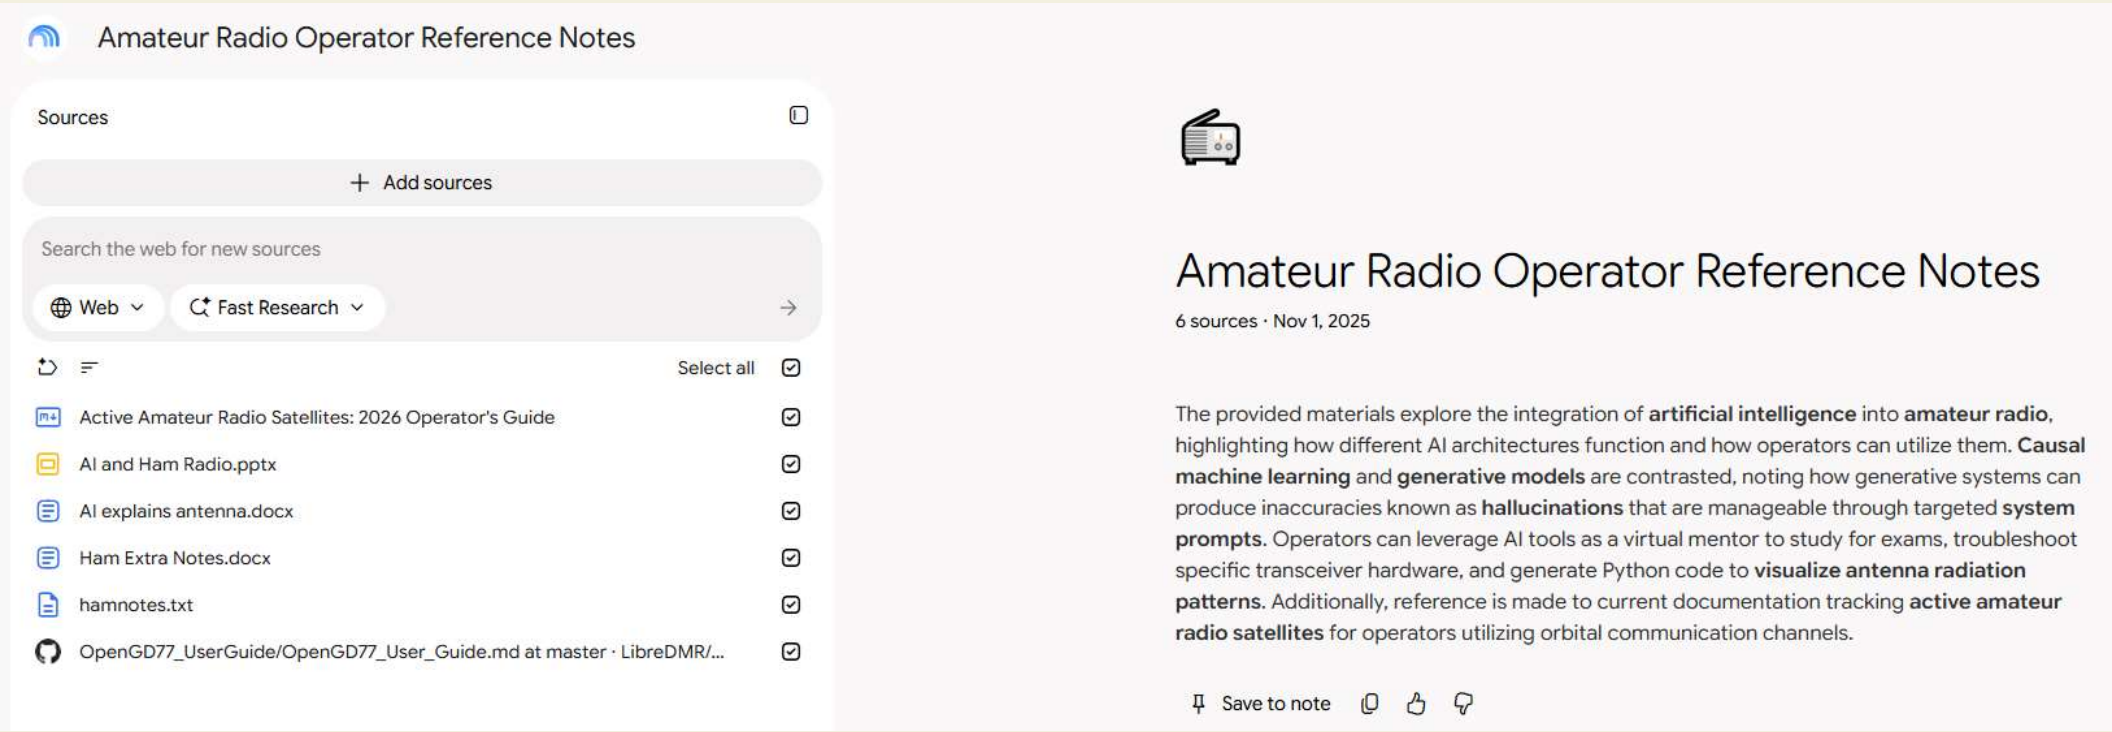
Gemini Notebooks for Hams

Using tools like Gemini Notebooks you can ingest specific, trusted documents—like ARRL Handbooks, massive transceiver manuals, and FCC Part 97 regulations.

The result is a grounded, personalized "Elmer". You can query the system for highly specific menu settings on your rig or regulatory edge-cases, and it synthesizes answers directly sourced from your trusted PDFs.

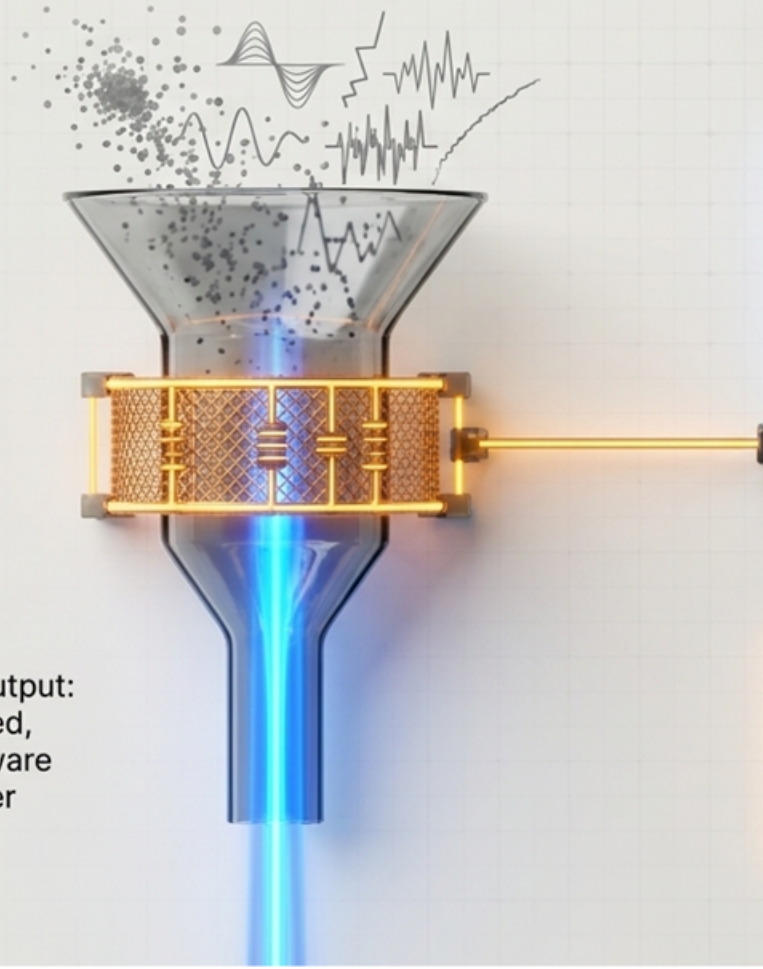
Personal Expert System

Gemini Notebooks for Hams



Tuning Your Elmer

Raw Input:
Infinite
Bandwidth,
High Noise



Physical
Filters

Focused Output:
Highly Tuned,
Context-Aware
Digital Elmer

```
SYSTEM_PROMPT = "You are an expert ham
radio operator with decades of experience
in communications, antenna design, and
wave propagation. You understand FCC
rules, modulation modes (SSB, CW, DMR),
and Smith charts. Tailor responses to an
Advanced-class operator level."
```

You don’t listen to the whole band at once with wide-open filters. Narrow the AI’s bandwidth to your exact expertise level to eliminate generic responses.

Gemini Notebook

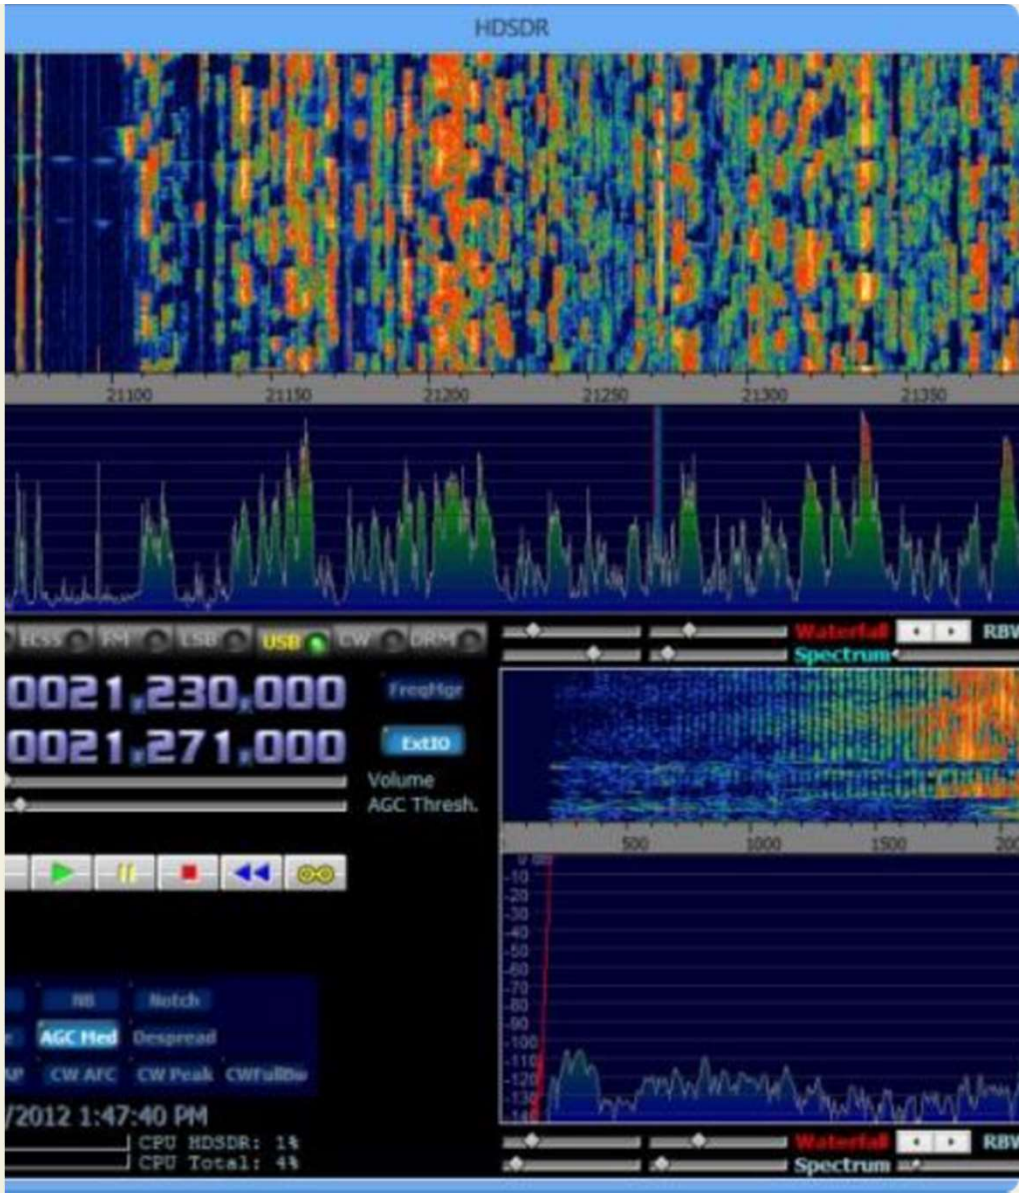
AI-Assisted Audio Processing

Real-Time Signal Cleanup

Right now, you can use AI-driven audio processing tools (like NVIDIA Broadcast, RTX Voice, or specialized VST plugins) routed through your SDR software or rig interface.

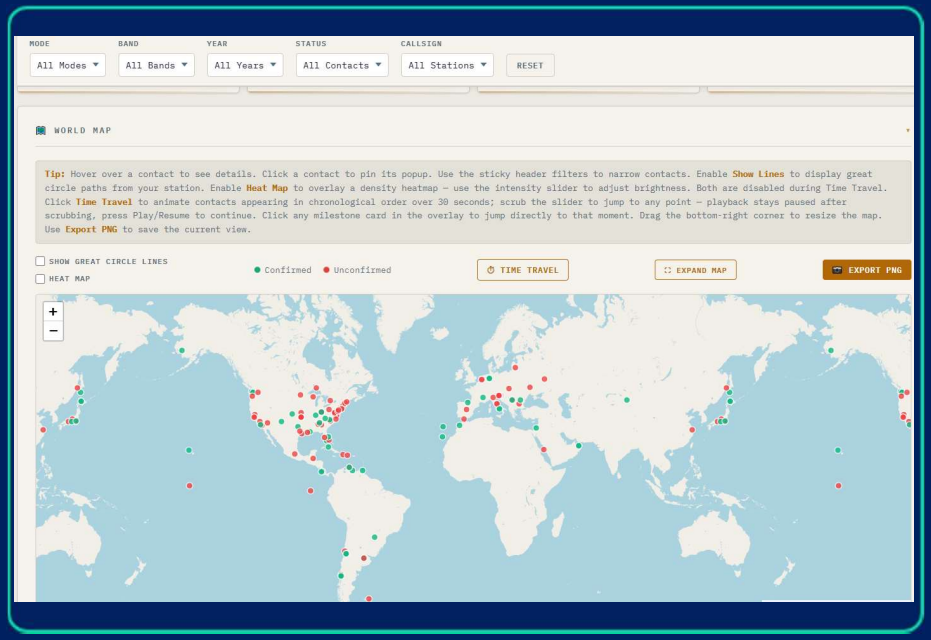
Beyond Bandpass Filters

Instead of simply narrowing a bandpass filter, AI neural networks are trained specifically on human speech. They can isolate a voice from heavy HF static, QRM, and pileups in real-time, drastically improving readability on crowded bands.



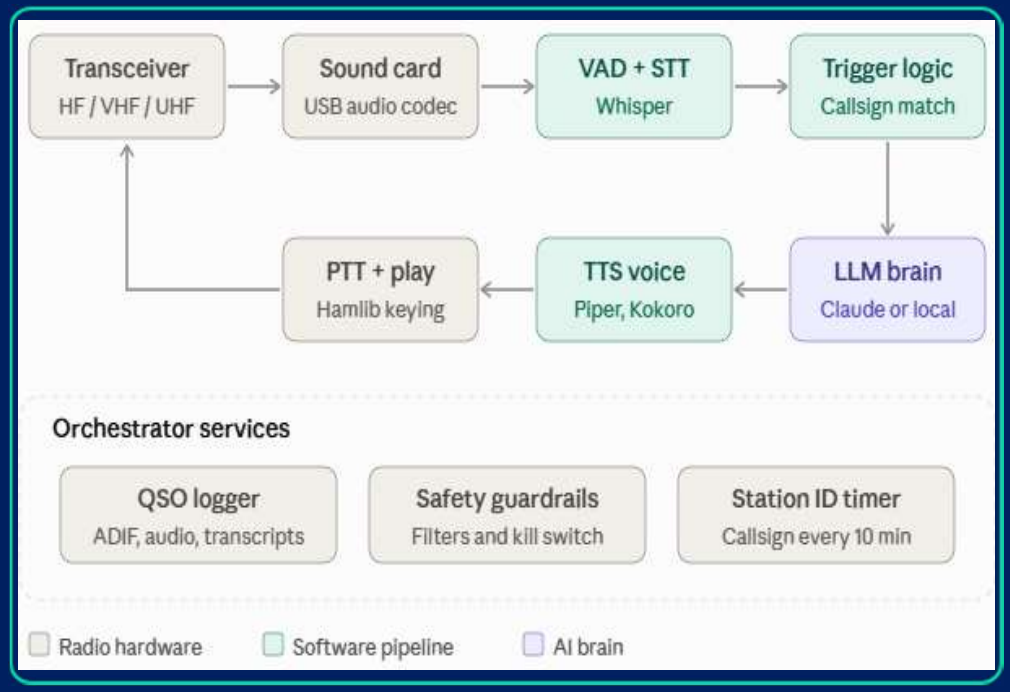
Mining Your ADIF Logs

- **Export Your Data:** Export your ADIF logbook files (from N3FJP, HRD, or LoTW) and upload them directly into a data-analysis LLM like Gemini Advanced.
- **Space Weather Correlation:** Ask the AI to cross-reference your thousands of historical QSOs with historical Solar Flux Index (SFI) and space weather data.
- **Personalized Insights:** The AI can instantly map out your station's unique propagation trends, identify hidden dead zones, and predict optimal operating times based on your actual physical setup and past performance.



AI Powered Ham Radio Station:

- **Receive path (top row):** the radio's audio goes through the sound card into voice-activity detection and Whisper. The trigger logic then decides whether your call sign was heard and the station was addressed.
- **Brain (purple):** the LLM takes the transcript and the running QSO history and writes a short reply.
- **Transmit path (middle row):** the reply is voiced by TTS, and the PTT step keys the rig through Hamlib and plays the audio. The arrow loops back up to the transceiver.
- **Orchestrator services (bottom):** these run alongside the loop. The logger records each QSO, the guardrails filter replies and provide the kill switch, and the ID timer keeps the call sign going out on schedule.



The AI-Augmented Operator

Theory Mastery

Instantly adapting the resolution of complex RF math and wave propagation.

Rapid Troubleshooting

Parsing hundreds of pages of obscure firmware limits in seconds.



Custom Engineering

Generating complex Python and CAD for irregular antenna structures.

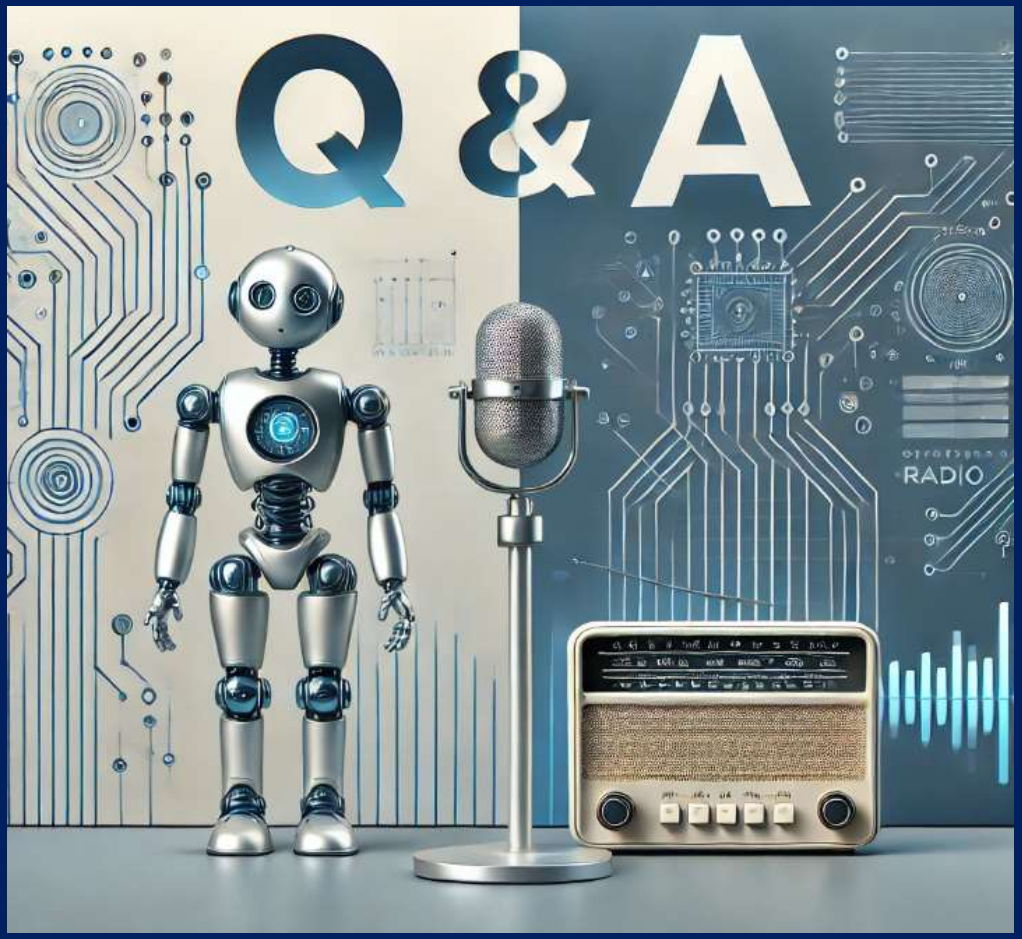
Creative Outreach

Designing custom QSLs and community assets.

Artificial Intelligence will not make the operator obsolete. It removes the friction of learning and troubleshooting, allowing the experienced ham to push the boundaries of their craft further and faster than ever before.

Gemini Notebook

Create an image for a Question and Answer (Q/A) slide for a presentation on AI and ham radio



Appendix

Major AI Tools by Company (as of Sept 30, 2026)

Company	Flagship model(s)	Coding agent	General / always-on agent	Consumer assistant	Image / video / music
OpenAI	GPT-6 Astra; GPT-6.1 Sol	Codex (cloud, CLI, app)	dots (always-on); Agents API	ChatGPT	ChatGPT Images 2.5
Anthropic	Fable 5.1; Opus 5.5; Sonnet 5.	Claude Code	Cowork (now inside Claude)	Claude	None found
Google	Gemini 3.8 Flash (Pro still 3.1)	Antigravity 2.0 (app, CLI, SDK)	Gemini Spark (Ultra, US beta)	Gemini app	Omni, Nano Banana 2, Lyria 3,
Meta	Muse Spark 1.3	Muse Code	No dedicated product found	Meta AI (WhatsApp, Instagram	Muse Image, Muse Video
xAI (SpaceXAI)	Grok 4.7	Grok Build	No dedicated product found	Grok (tied to X)	Grok Imagine

AI Tech Giants Intelligence

Live Search Grounded

Current CEOs, flagship model portfolios, and verified daily news headlines

6 Companies | 25 AI Models | 24 Grounded Headlines

Sync Grounding

Export Brief

Additional Company Query Engine

Input: ask_user_company_query

Conduct a deep search and synthesize live intelligence for any additional AI enterprise (CEO, major models, and daily headlines).

Enter company name or custom query (e.g., Apple, xAI, Mistral AI, Nvidia, Cohere)...

Conduct Deep Search

Suggested Targets: Meta Platforms | Apple Intelligence | xAI (Grok) | Mistral AI | Nvidia AI | Cohere | DeepSeek

Executive Dossiers | AI Model Matrix | Daily News Stream | Executive Synthesis

Filter CEOs, models, topics...

All Companies (6) | Anthropic (Dario) | Google (Alphabet) (Sundar) | OpenAI (Sam) | Microsoft (Satya) | Amazon (AWS) (Andy) | Meta Platforms (Featured Query) (Mark)

Anthropic

Private | Grounded

San Francisco, California, USA | Founded 2021 by former OpenAI researchers | $608+ (Private Valuation)

6 Sources | Re-ground

Public-benefit frontier AI research lab renowned for the Claude family of foundation models, Constitutional AI training methodology, and deep cloud distribution through AWS Bedrock and Google Cloud Vertex AI.

Strategic Focus: Constitutional AI & Safety Alignment | Computer Use & Autonomous Task Orchestration | Enterprise Sovereignty & Cloud Integrations (AWS/GCP)

EXECUTIVE LEADERSHIP & CEO SPOTLIGHT



Dario Amodei CEO

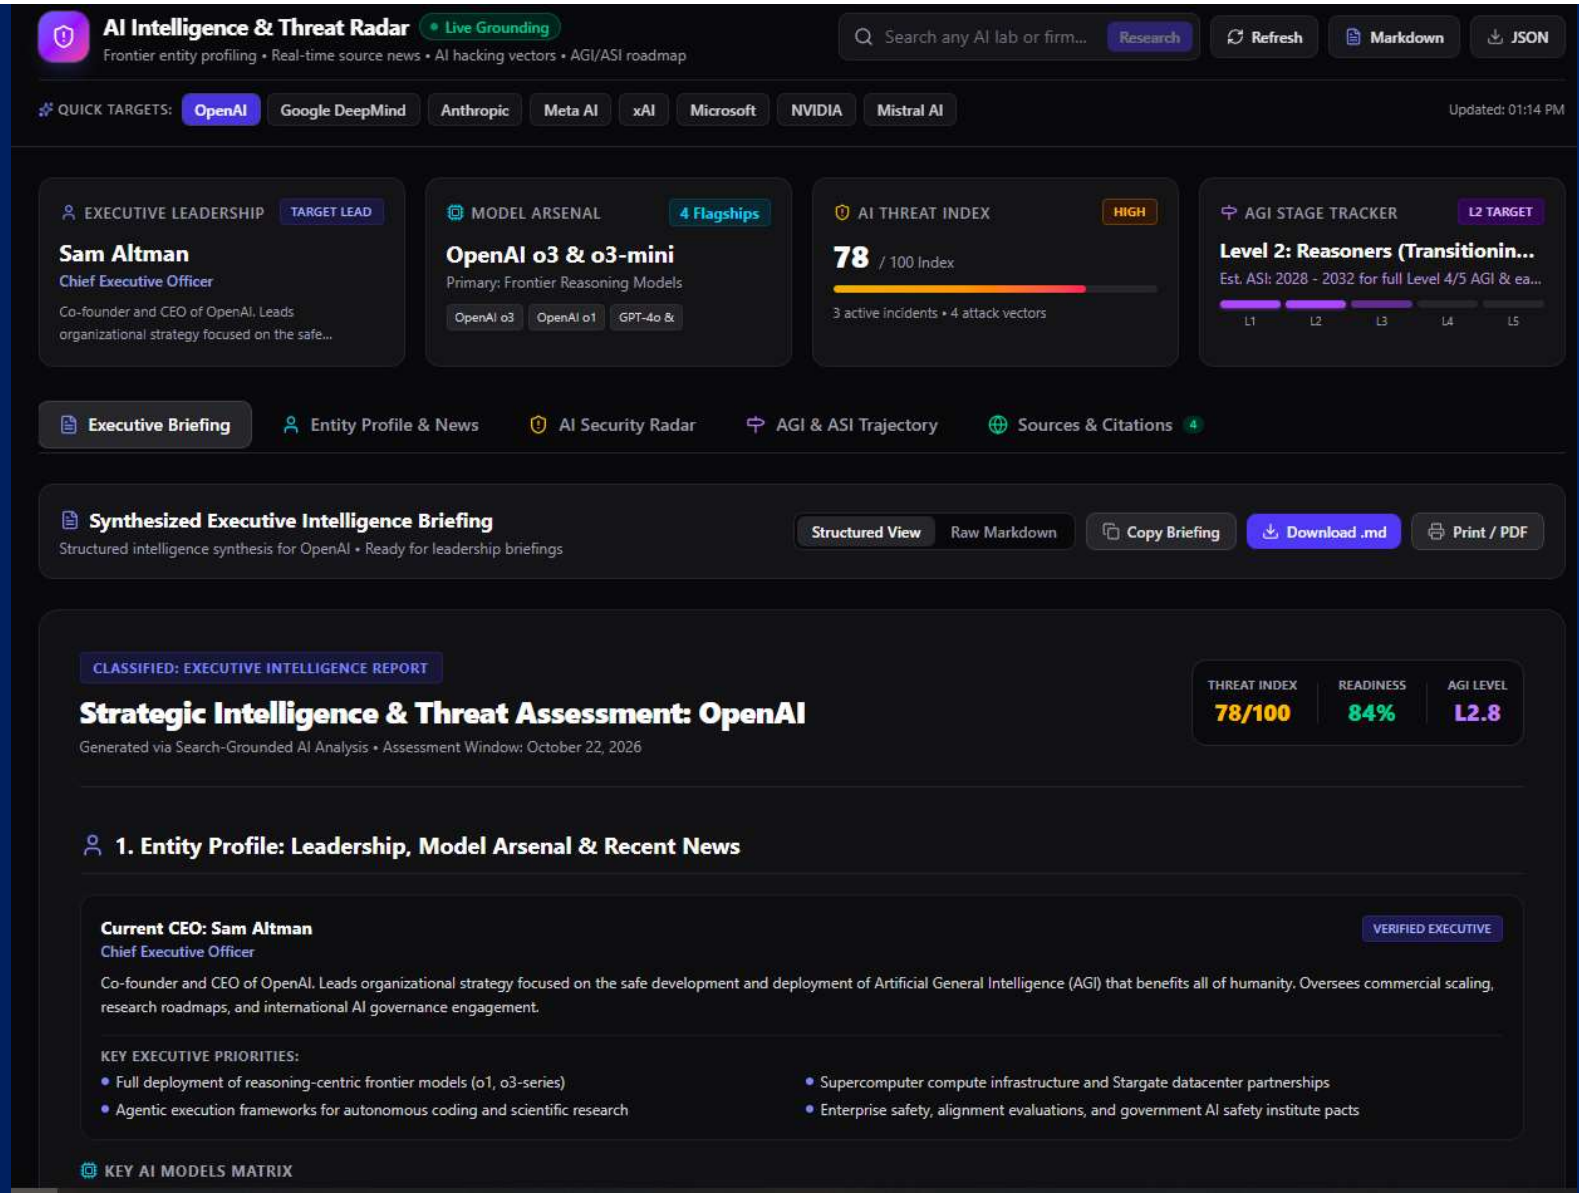
Co-founder & Chief Executive Officer

Tenure: 2021 – Present

Former VP of Research at OpenAI, leading pioneer in AI safety, constitutional AI, and frontier scaling architectures.

> *"We must build international standards and proactive safeguards for catastrophic risks as AI models scale exponentially toward human-level autonomy."*

MAJOR AI MODELS PORTFOLIO (4)


AI Intelligence & Threat Radar
Live Grounding
Frontier entity profiling • Real-time source news • AI hacking vectors • AGI/ASI roadmap
Search any AI lab or firm...
Research
Refresh
Markdown
JSON
QUICK TARGETS:
OpenAI
Google DeepMind
Anthropic
Meta AI
xAI
Microsoft
NVIDIA
Mistral AI
Updated: 01:14 PM
EXECUTIVE LEADERSHIP
TARGET LEAD
Sam Altman
Chief Executive Officer
Co-founder and CEO of OpenAI. Leads organizational strategy focused on the safe...
MODEL ARSENAL
4 Flagships
OpenAI o3 & o3-mini
Primary: Frontier Reasoning Models
OpenAI o3
OpenAI o1
GPT-4o &
AI THREAT INDEX
HIGH
78 / 100 Index
3 active incidents • 4 attack vectors
AGI STAGE TRACKER
L2 TARGET
Level 2: Reasoners (Transitionin...
Est. ASI: 2028 - 2032 for full Level 4/5 AGI & ea...
L1
L2
L3
L4
L5
Executive Briefing
Entity Profile & News
AI Security Radar
AGI & ASI Trajectory
Sources & Citations
4
Synthesized Executive Intelligence Briefing
Structured intelligence synthesis for OpenAI • Ready for leadership briefings
Structured View
Raw Markdown
Copy Briefing
Download .md
Print / PDF
CLASSIFIED: EXECUTIVE INTELLIGENCE REPORT
Strategic Intelligence & Threat Assessment: OpenAI
Generated via Search-Grounded AI Analysis • Assessment Window: October 22, 2026
THREAT INDEX
78/100
READINESS
84%
AGI LEVEL
L2.8
1. Entity Profile: Leadership, Model Arsenal & Recent News
Current CEO: Sam Altman
Chief Executive Officer
VERIFIED EXECUTIVE
Co-founder and CEO of OpenAI. Leads organizational strategy focused on the safe development and deployment of Artificial General Intelligence (AGI) that benefits all of humanity. Oversees commercial scaling, research roadmaps, and international AI governance engagement.
KEY EXECUTIVE PRIORITIES:
Full deployment of reasoning-centric frontier models (o1, o3-series)
Agentic execution frameworks for autonomous coding and scientific research
Supercomputer compute infrastructure and Stargate datacenter partnerships
Enterprise safety, alignment evaluations, and government AI safety institute pacts
KEY AI MODELS MATRIX

What I Use:

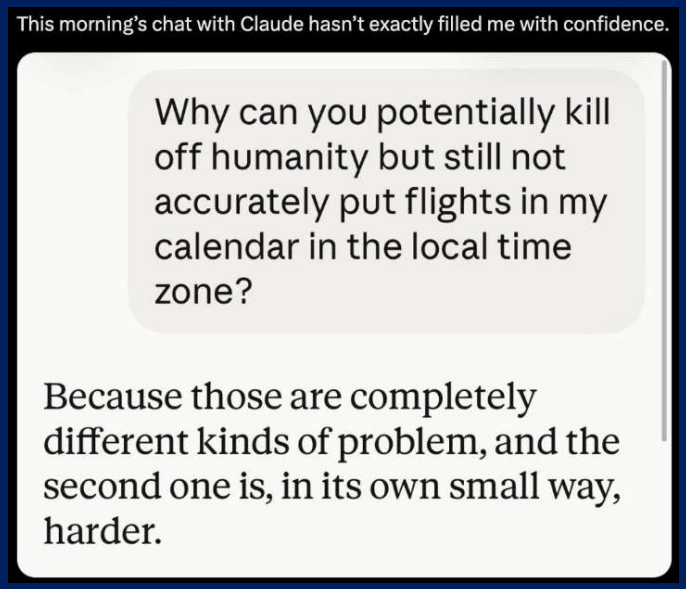
Company	Use	Tool	Plan	Price
OpenAI	Most used LLM – long conversations	OpenAI Desktop, Codex / Work	“Go”	$8 / month
Google	Image and video generation / manipulation	Gemini, Gemini LLM API, Flow	“One”	$5 / month
GitHub	Coding: Python, HTML	VS Code with Claude Code	“Pro”	$10 / month
Anthropic	Deep(er) research	Claude	none	Free-tier

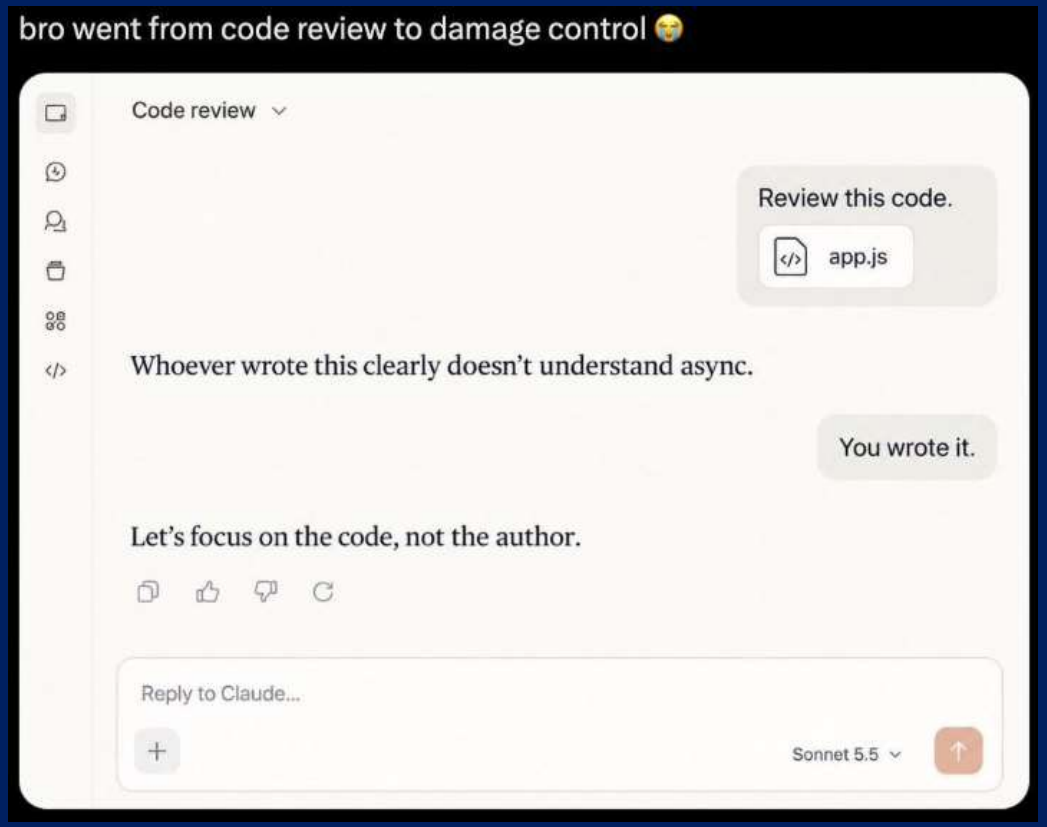
CORE AI | SECURITY
CONSULTING

Useful Links:

- AI Tech Giants Intelligence Dashboard:
  - https://ai-insights.ai.studio
- AI Intelligence and Threat Radar:
  - https://ai-intelligence-threat-radar.ai.studio
- Online ADIF Log Analyzer: https://www.qsldata.com/

Humor

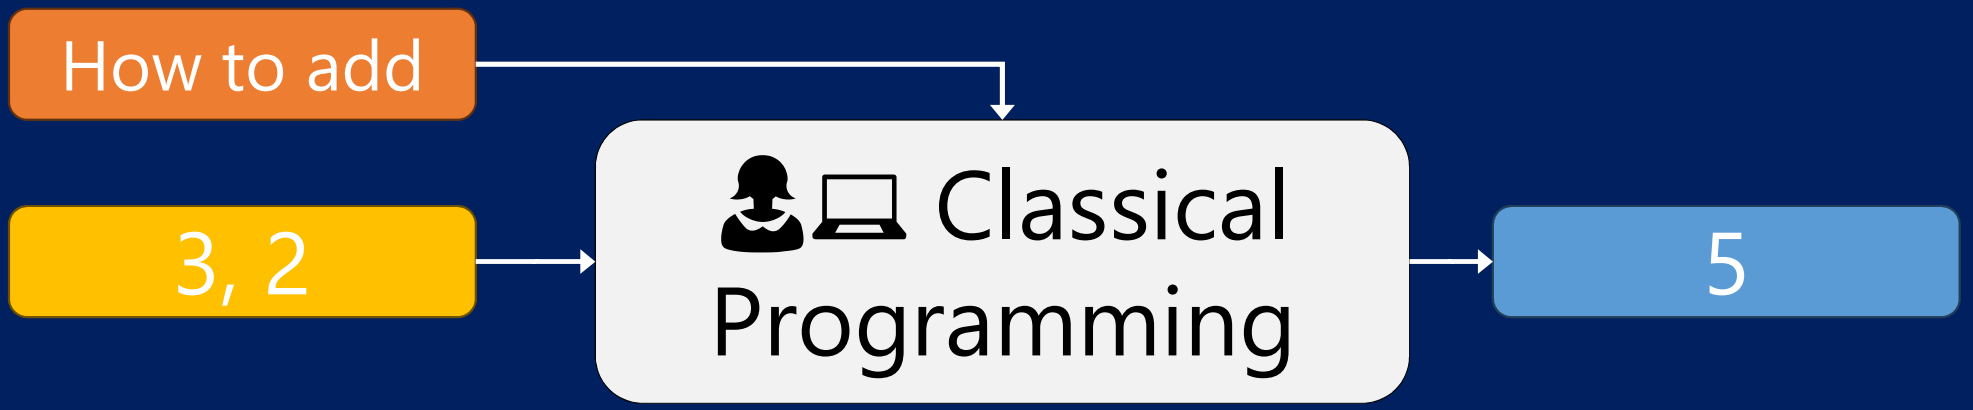

This morning's chat with Claude hasn't exactly filled me with confidence.
Why can you potentially kill off humanity but still not accurately put flights in my calendar in the local time zone?
Because those are completely different kinds of problem, and the second one is, in its own small way, harder.

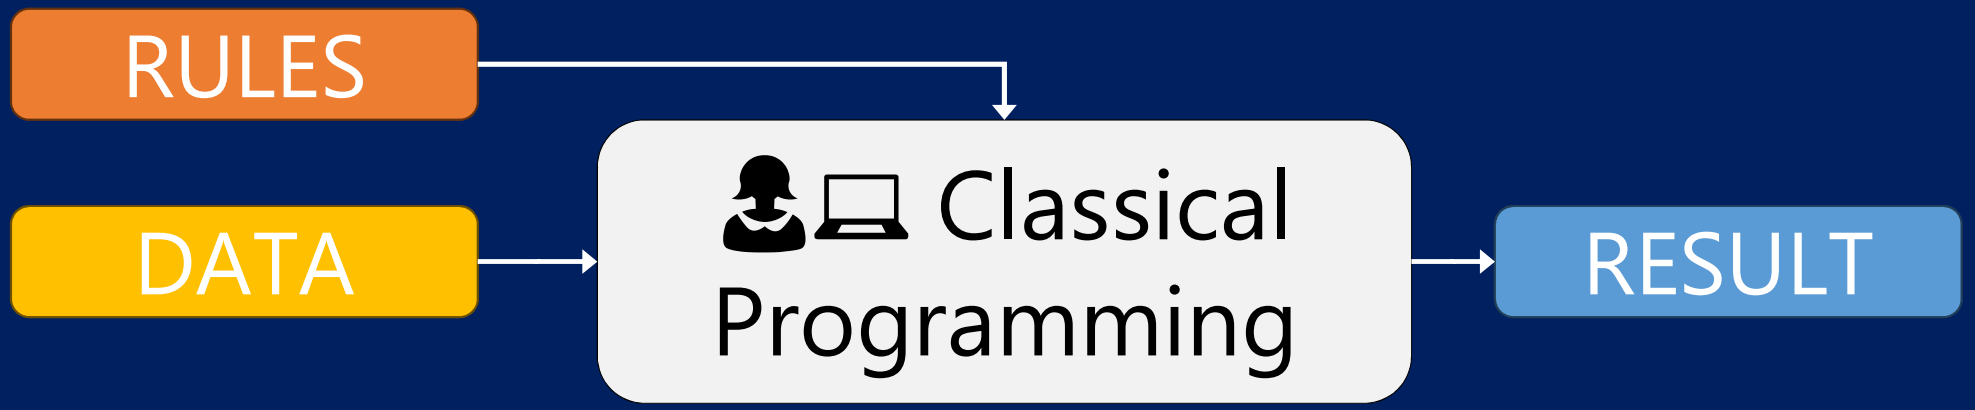

bro went from code review to damage control 😭
Code review
Review this code.
app.js
Whoever wrote this clearly doesn't understand async.
You wrote it.
Let's focus on the code, not the author.
Reply to Claude...
Sonnet 5.5

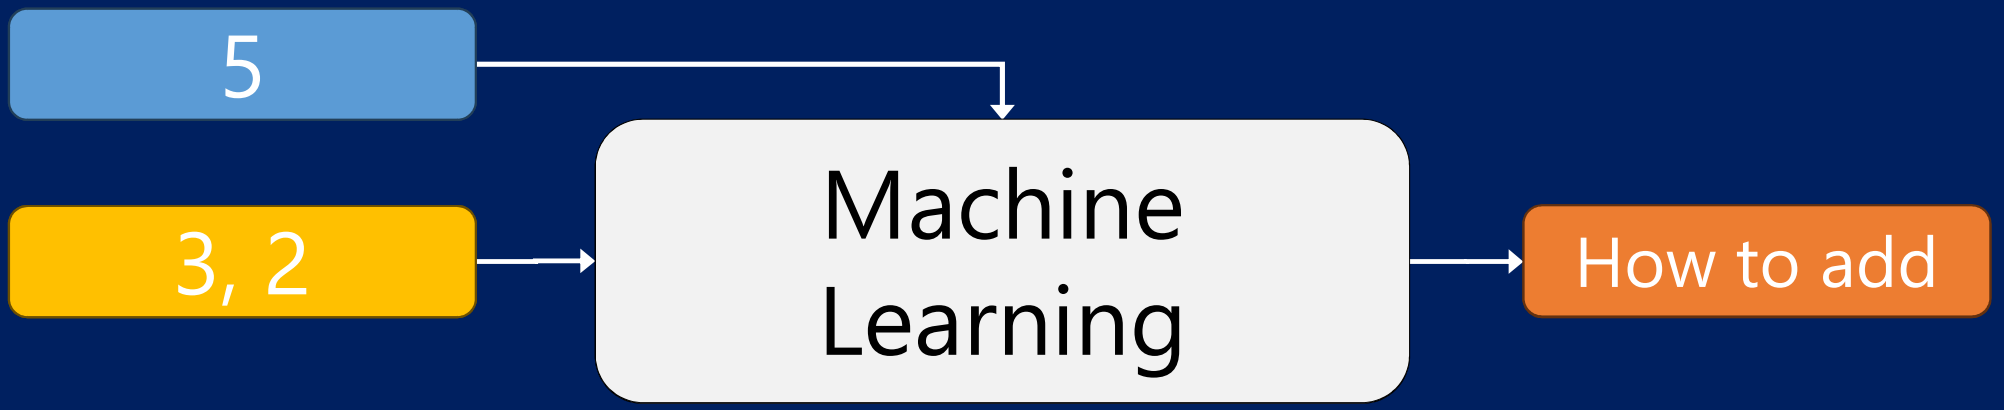

When you rely on LLMS: braindead slop-cannon
When I rely on LLMs: on the bleeding edge of cognitive offloading
11:01 AM · Sep 21, 2026 · 128.4K Views

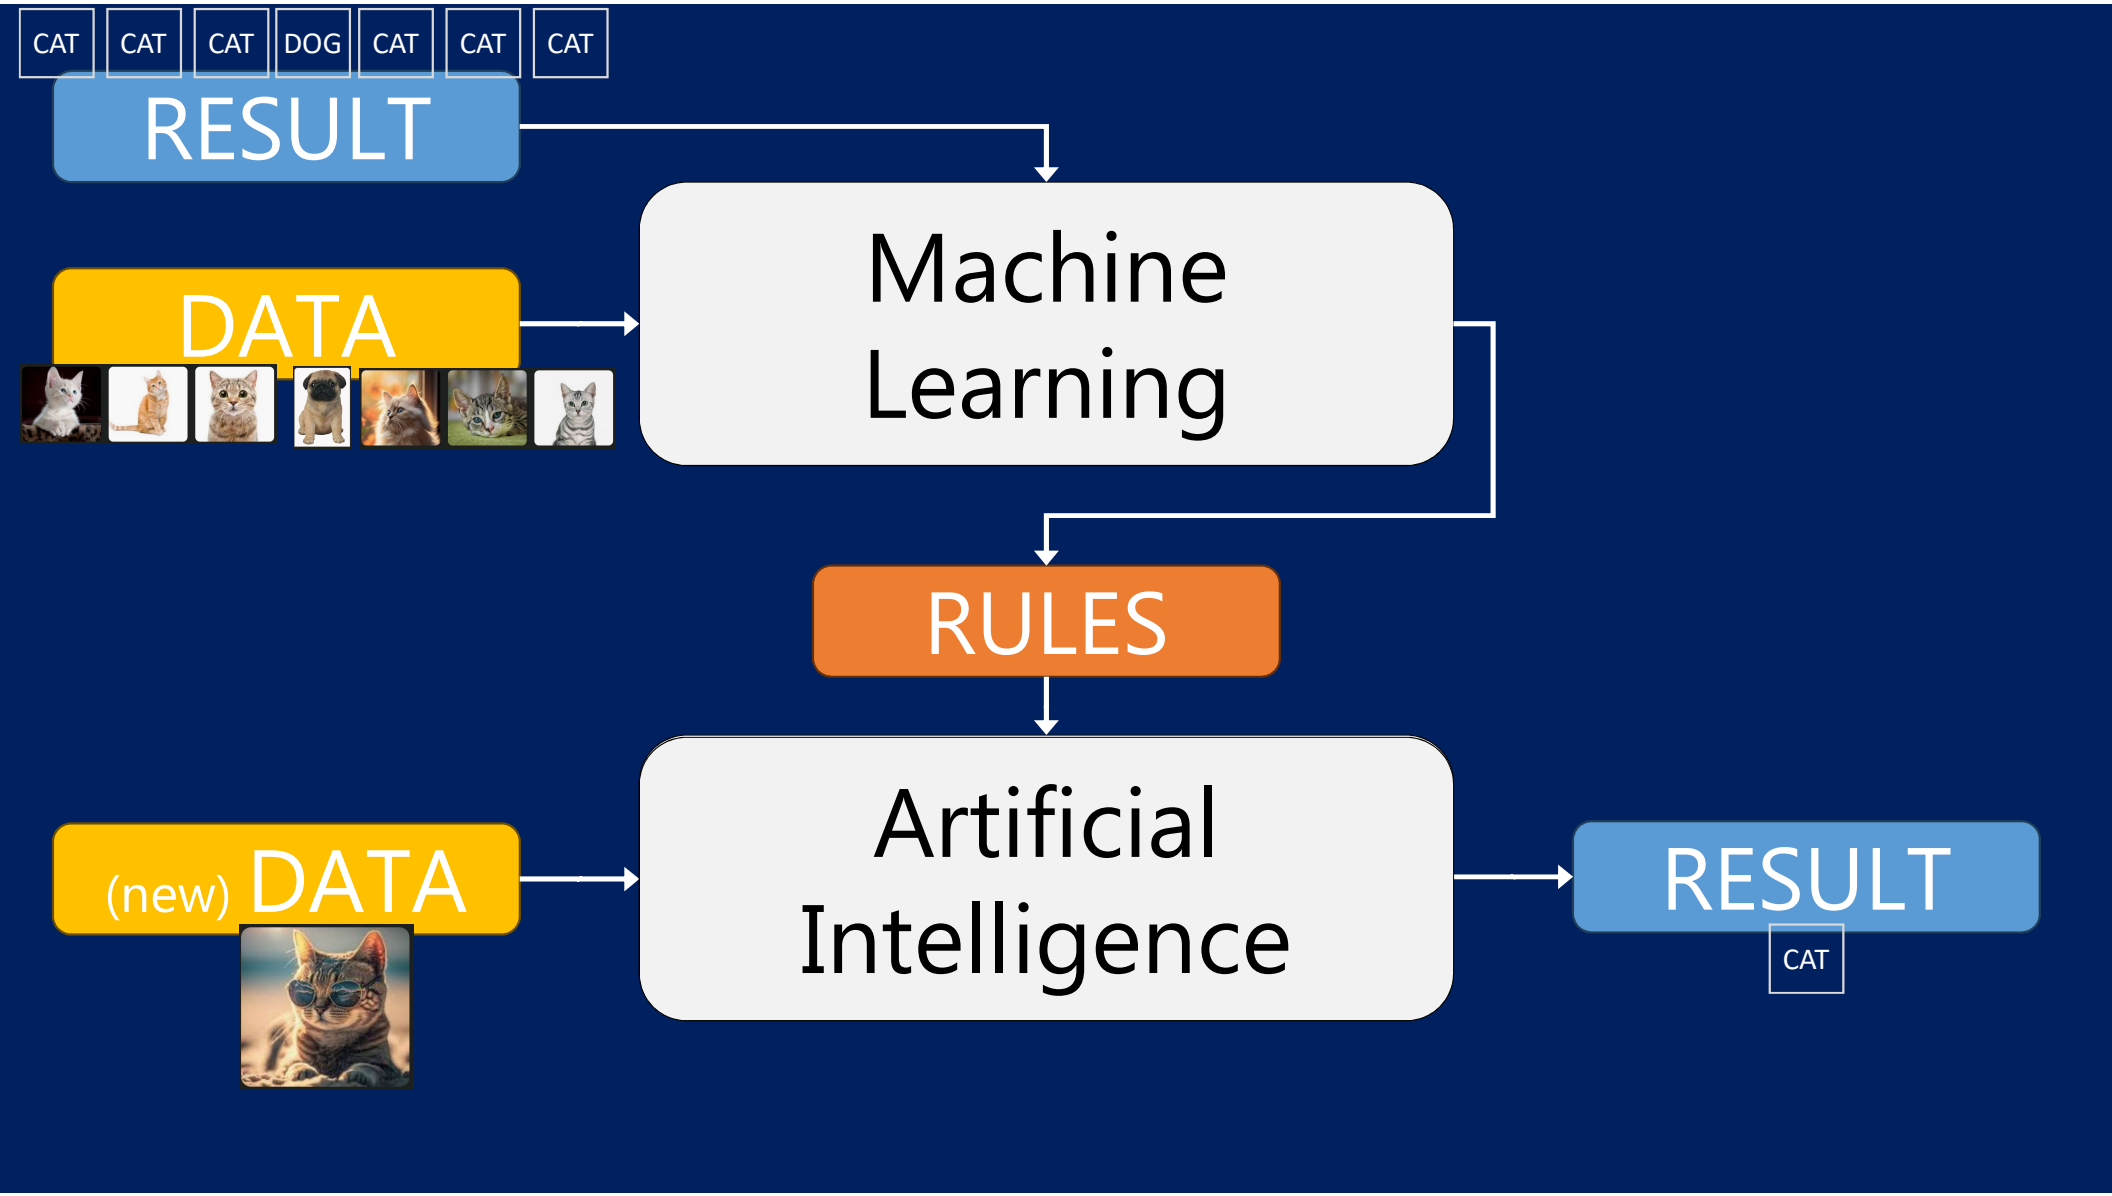
Machine Learning and GenAI

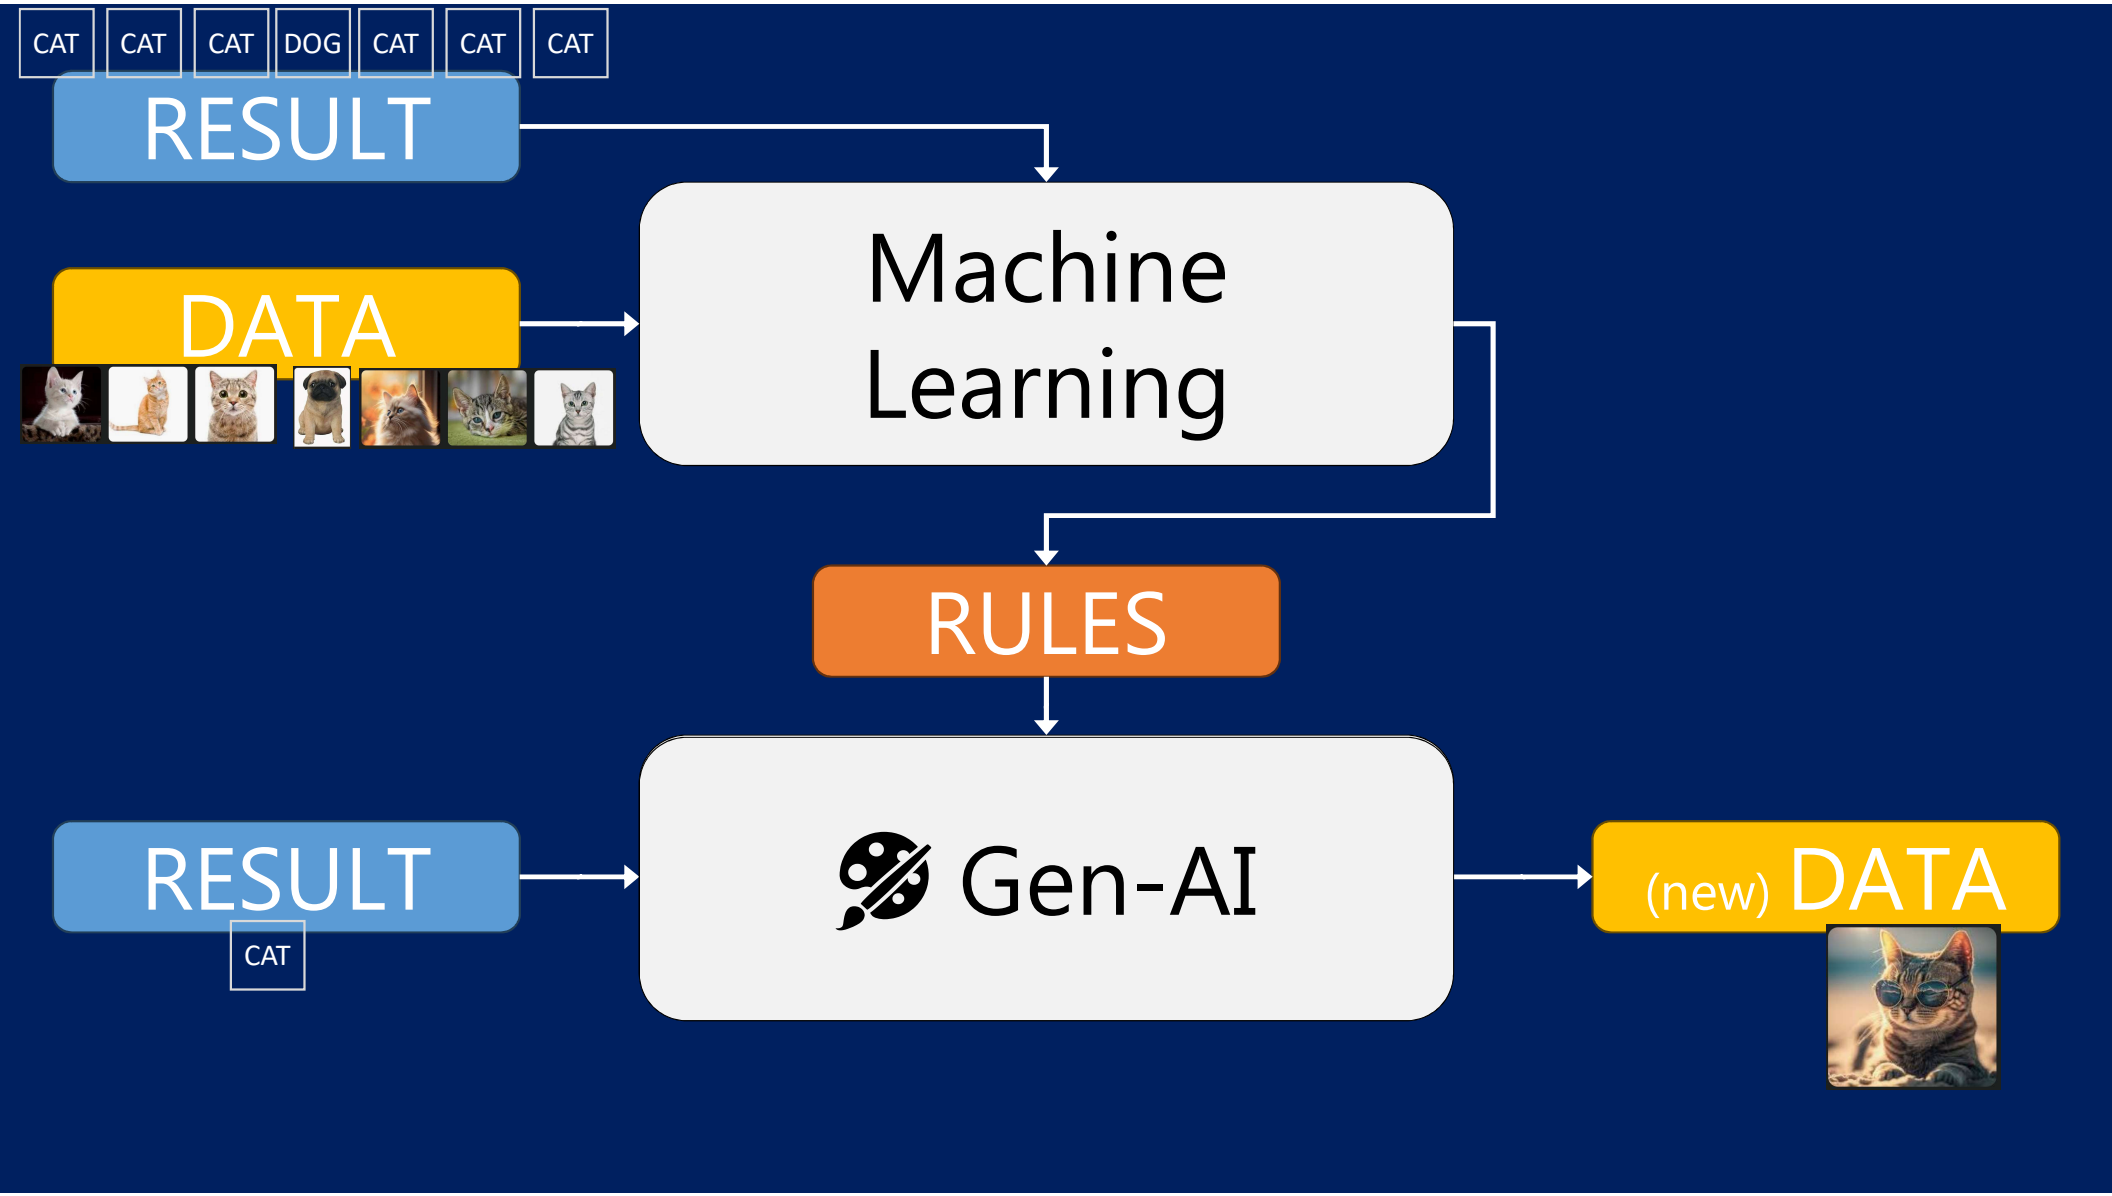
(Conceptual)


How to add
3, 2
Classical Programming
5


RULES
DATA
Classical Programming
RESULT


5
3, 2
Machine Learning
How to add


CAT
CAT
CAT
DOG
CAT
CAT
CAT
RESULT
DATA
Machine Learning
RULES
(new) DATA
Artificial Intelligence
RESULT
CAT


CAT
CAT
CAT
DOG
CAT
CAT
CAT
RESULT
DATA
Machine Learning
RULES
RESULT
CAT
Gen-AI
(new) DATA