AI Fundamentals for Govt. Leaders

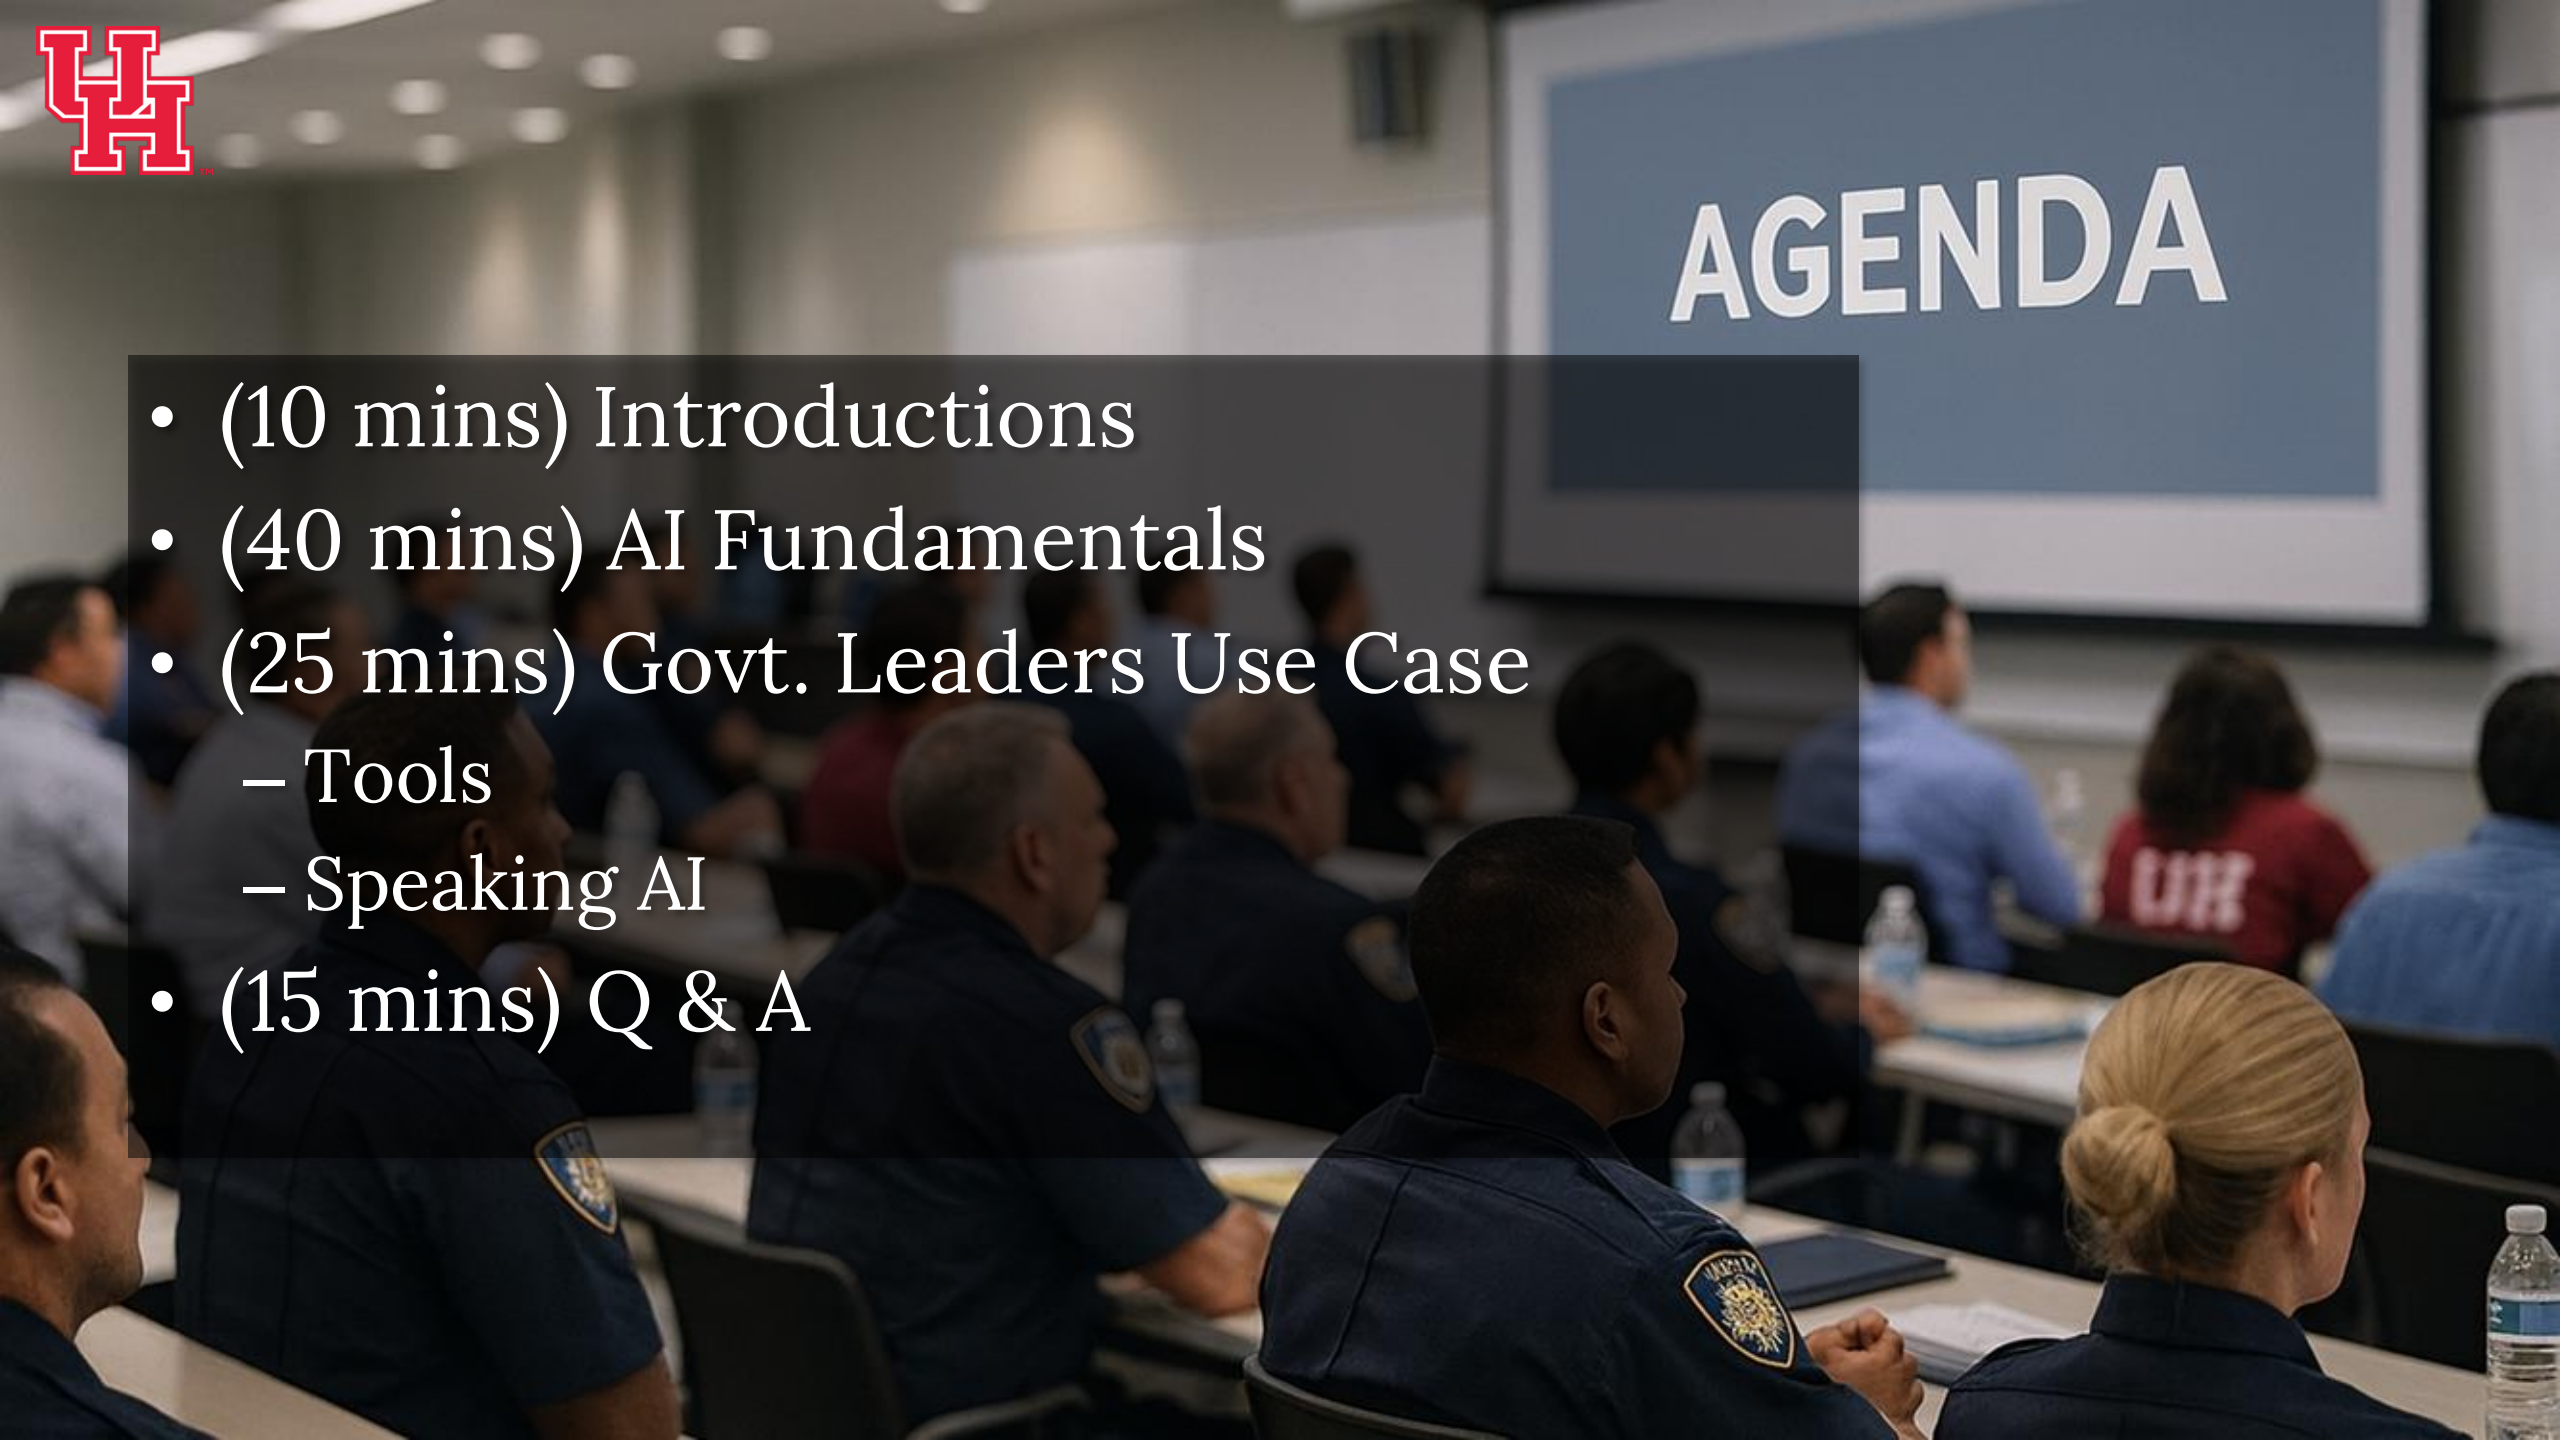

- (10 mins) Introductions
- (40 mins) AI Fundamentals
- (25 mins) Govt. Leaders Use Case
  - Tools
  - Speaking AI
- (15 mins) Q & A

Coca-Cola

lifelong learner – 3x

AI Fundamentals for Govt. Leaders

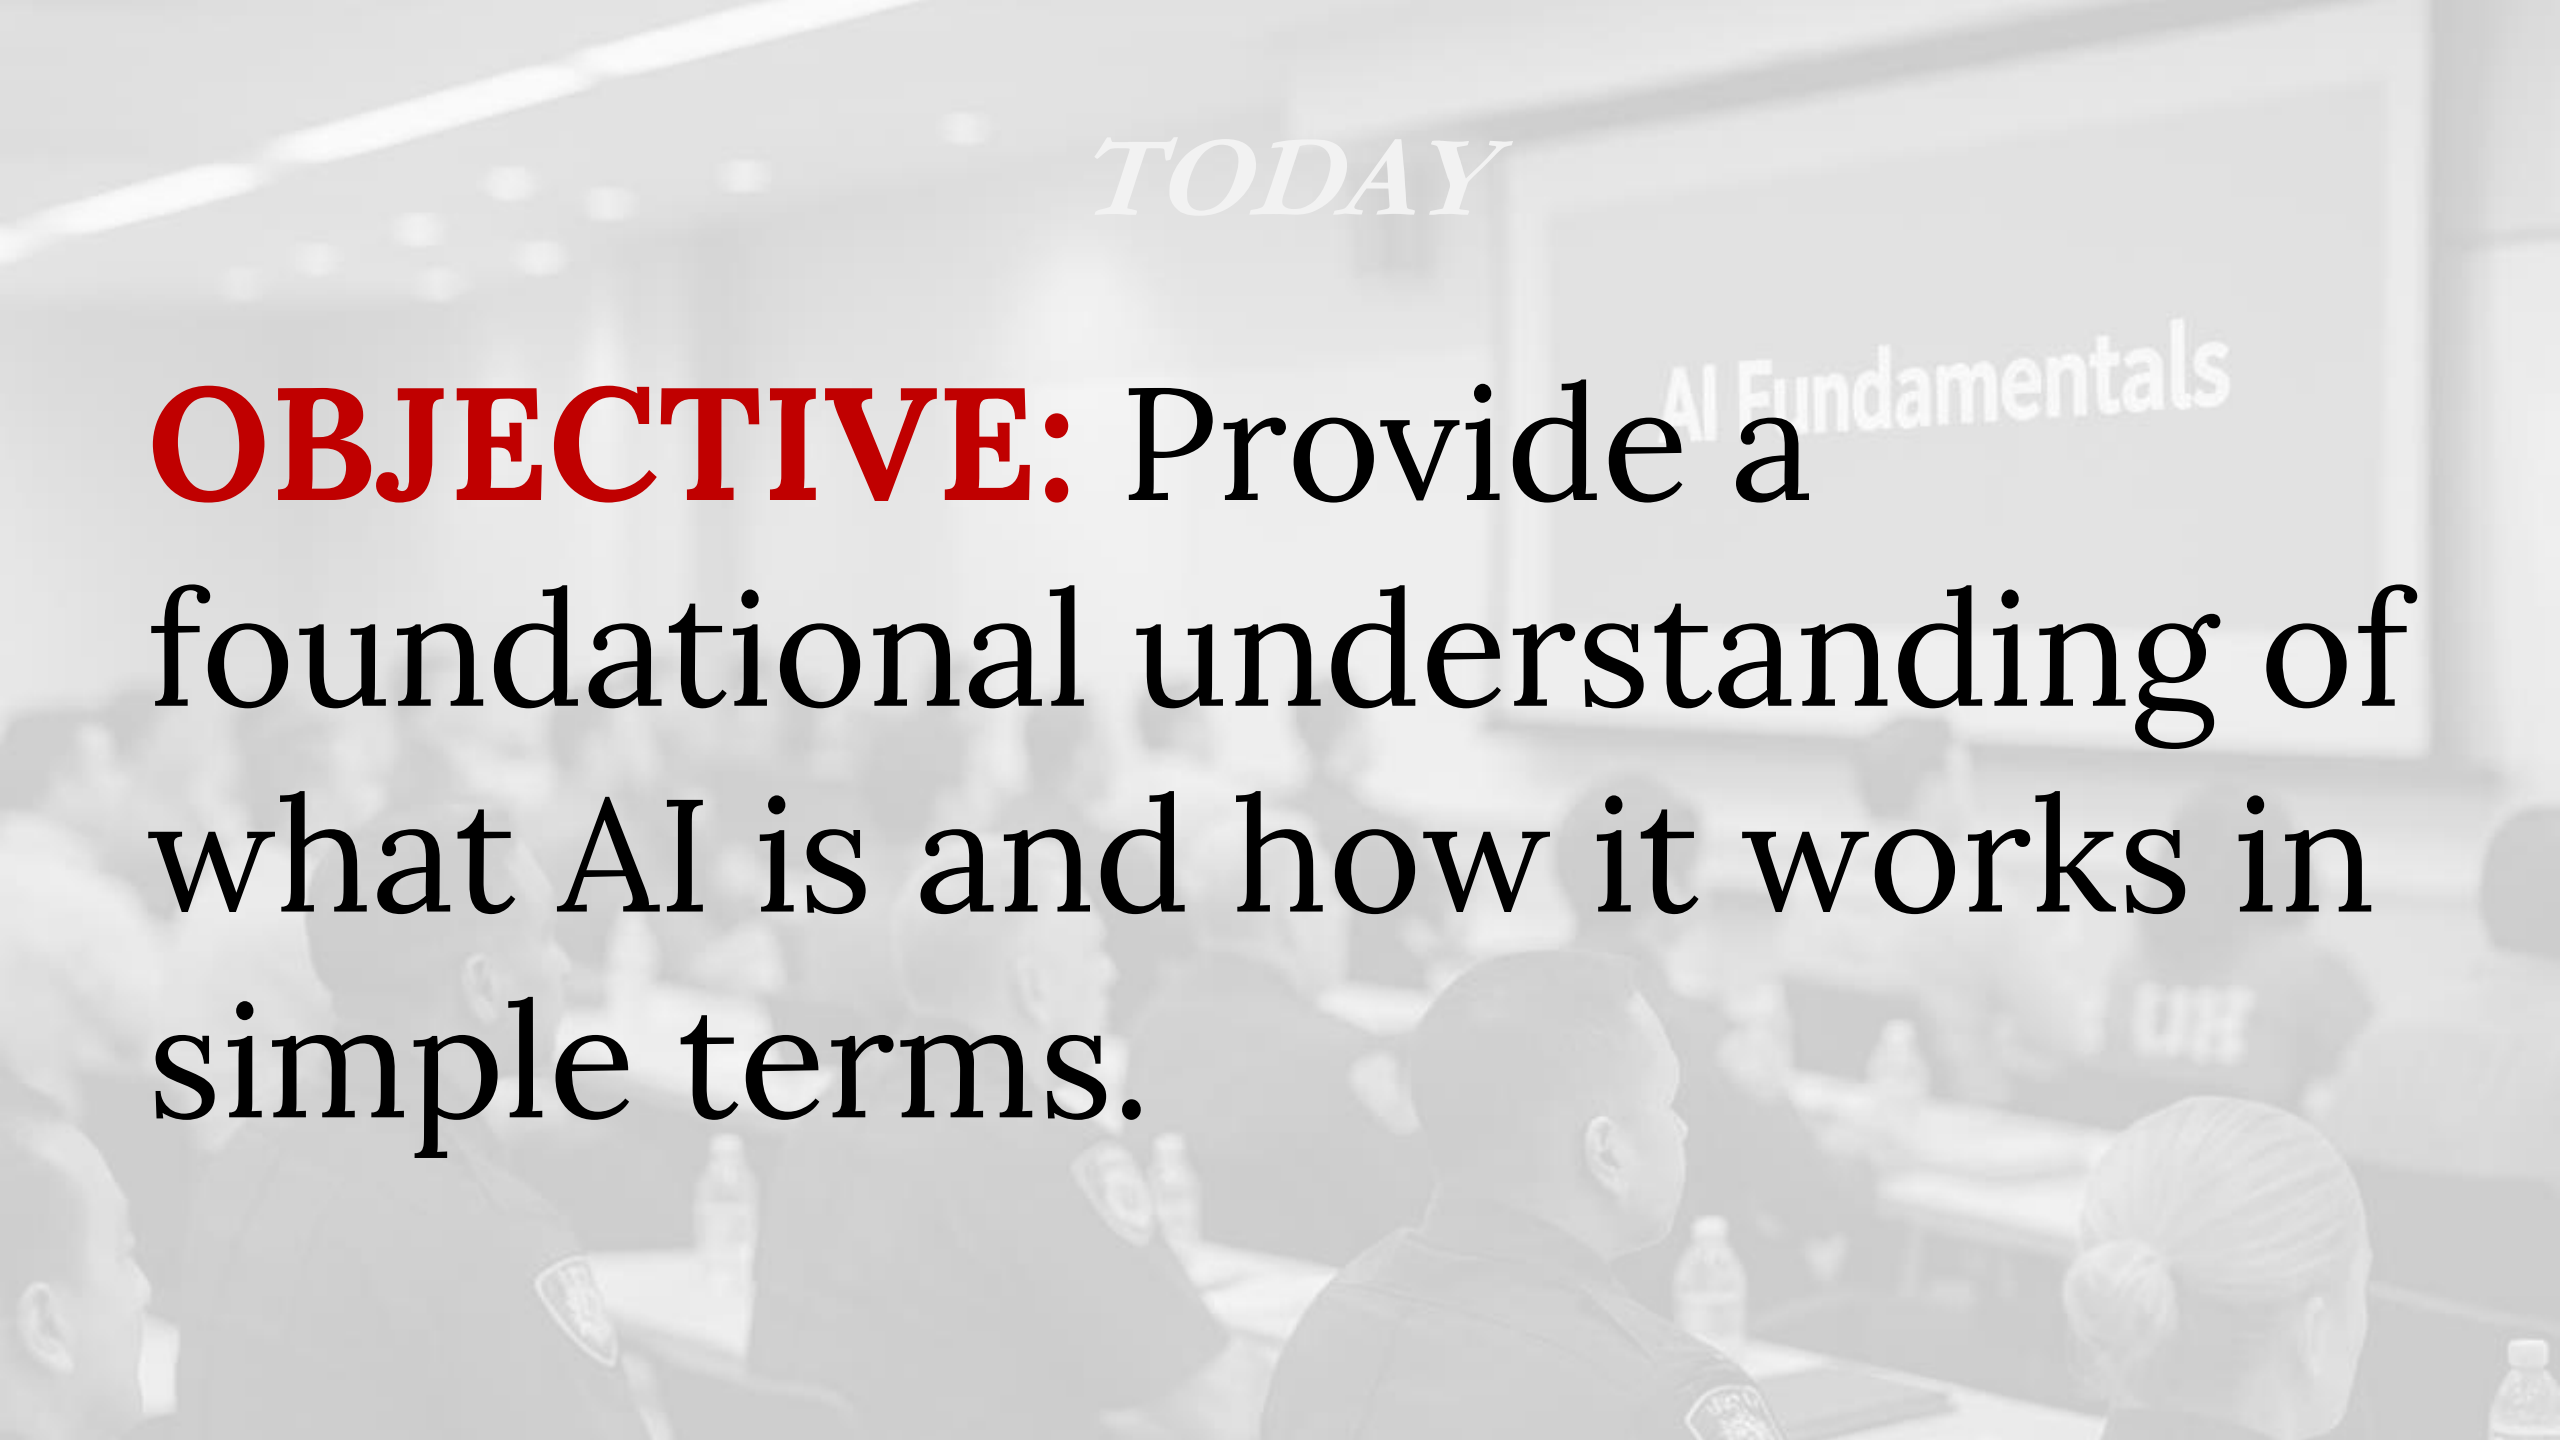

OBJECTIVE: Provide a foundational understanding of what AI is and how it works in simple terms.

Artificial Intelligence

AI is the ability of a machine to *mimic human intelligence*, such as learning from data, recognizing patterns, and making decisions.

WHAT WE IMAGINE AI HAS BECOME…

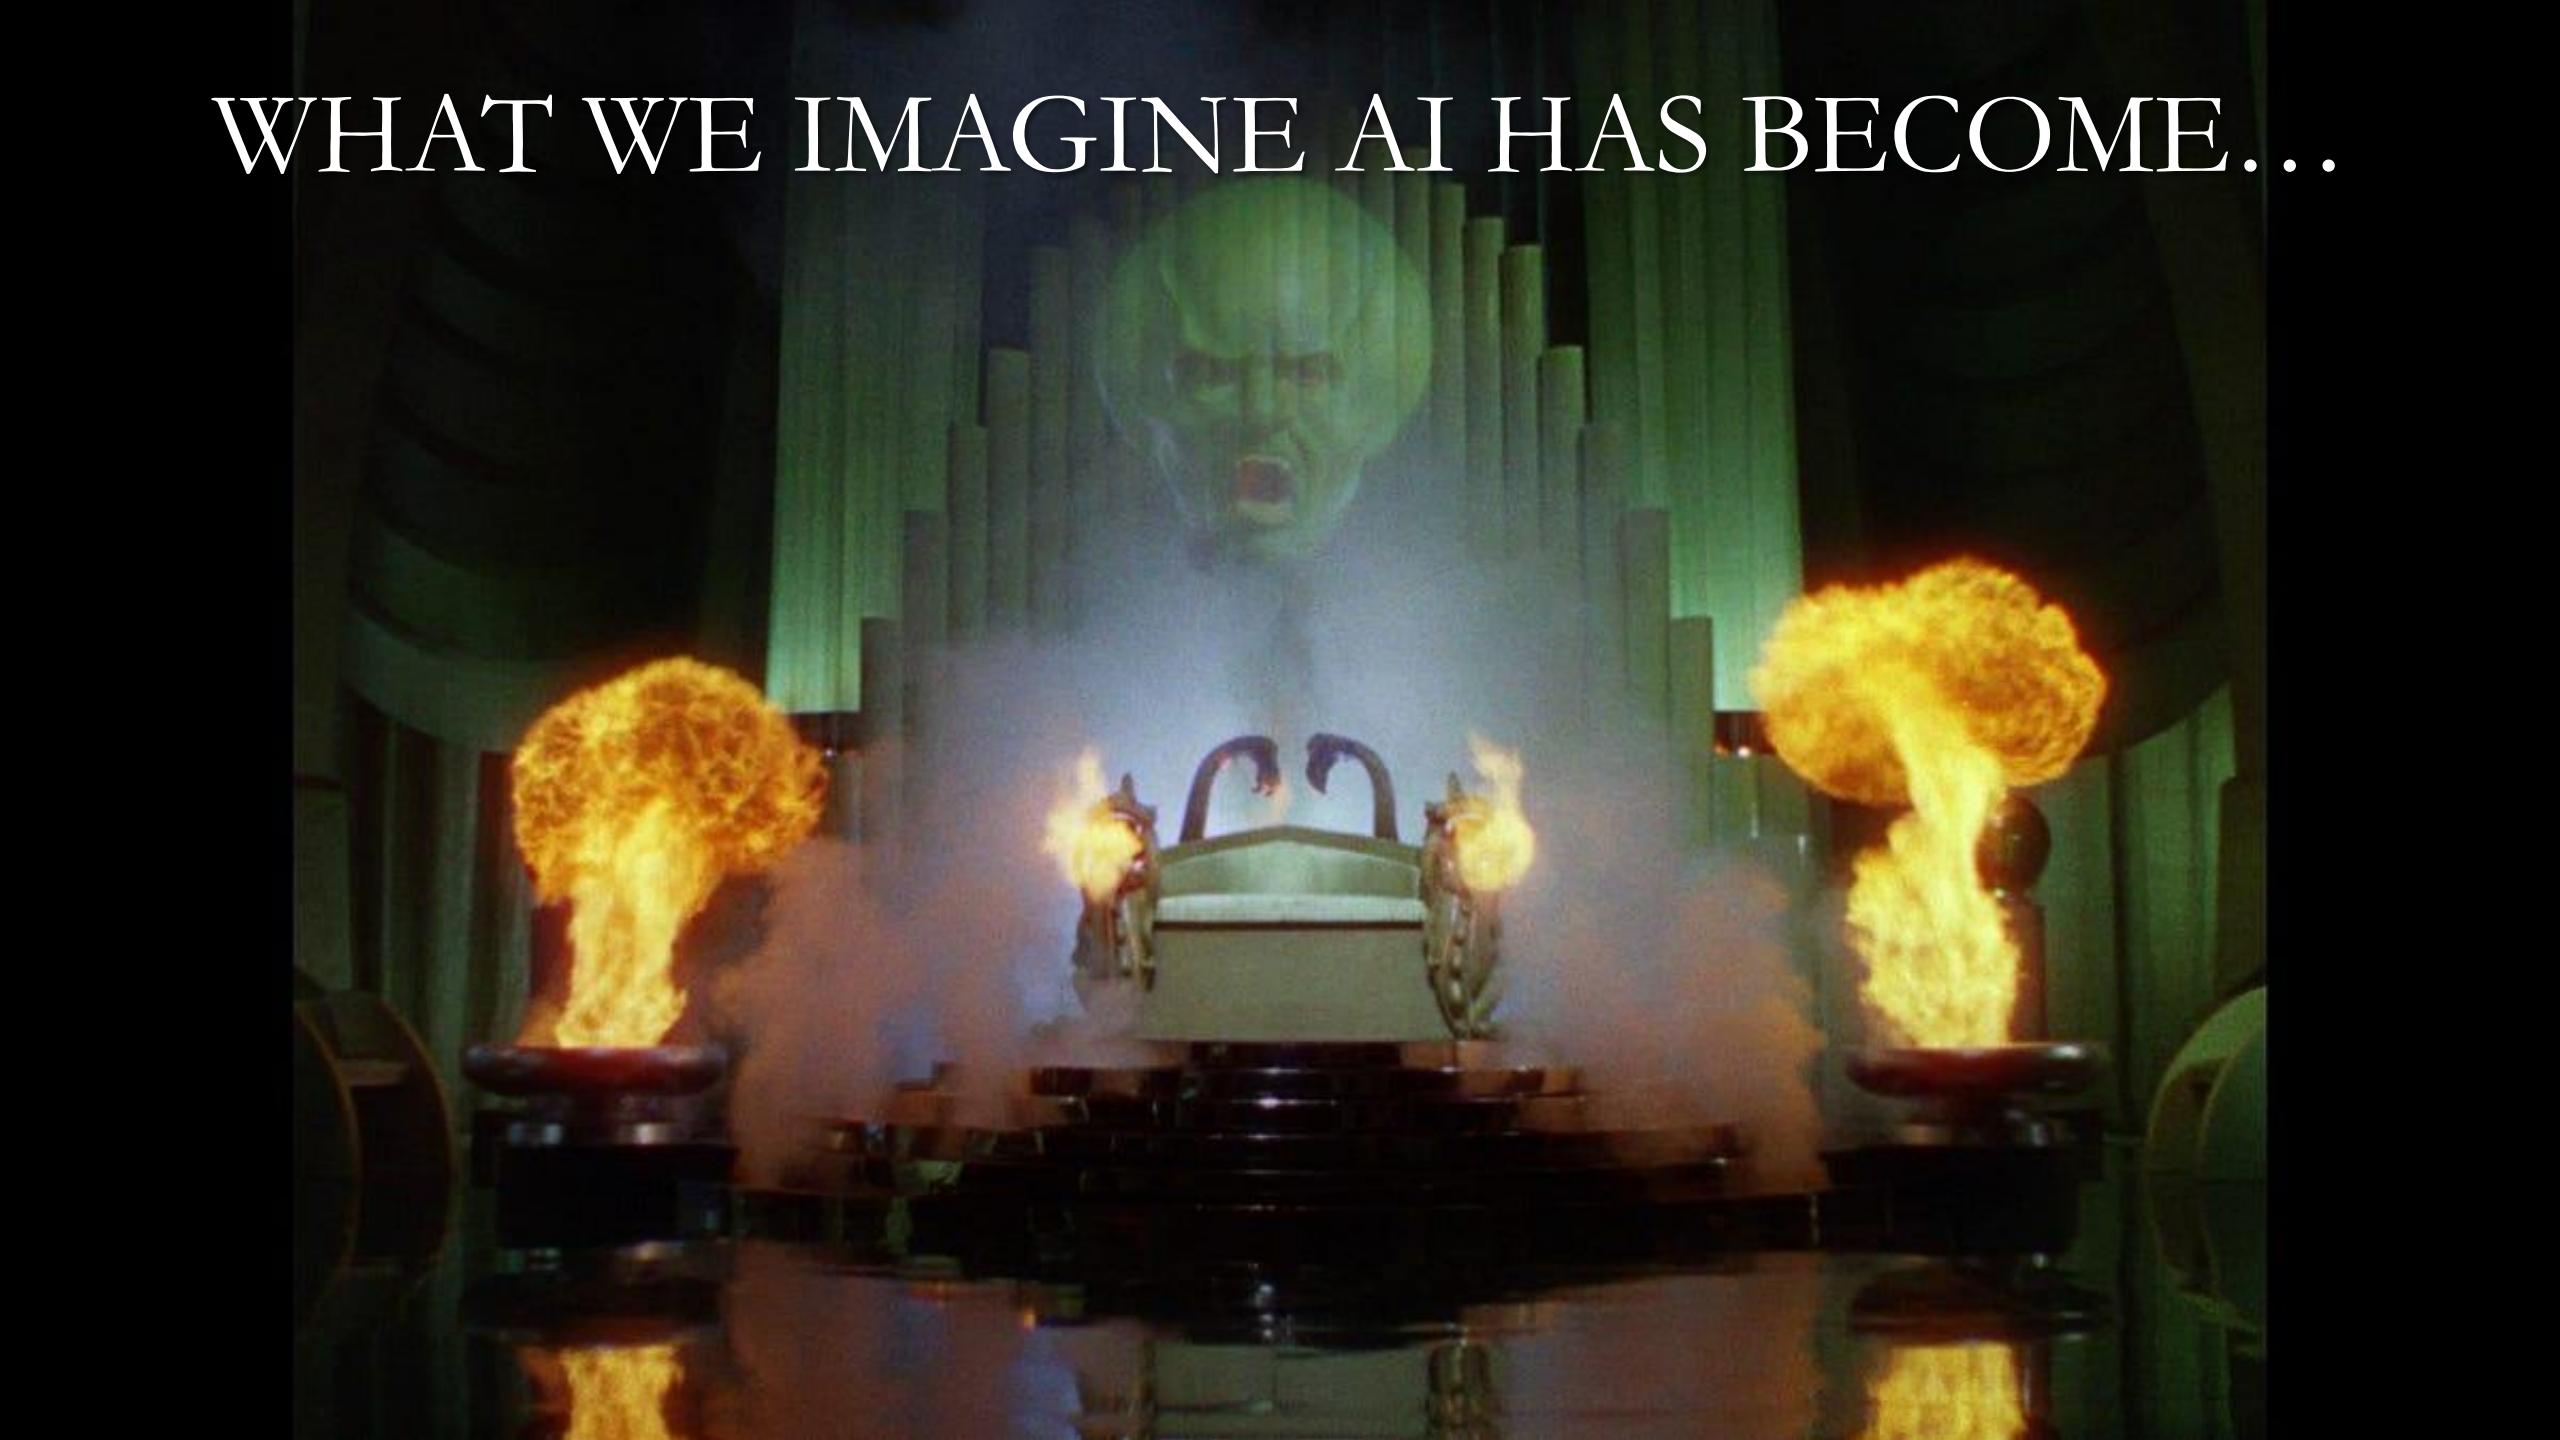


AS OF NOVEMBER 20, 2025

AI

MODELS EXPLAINED

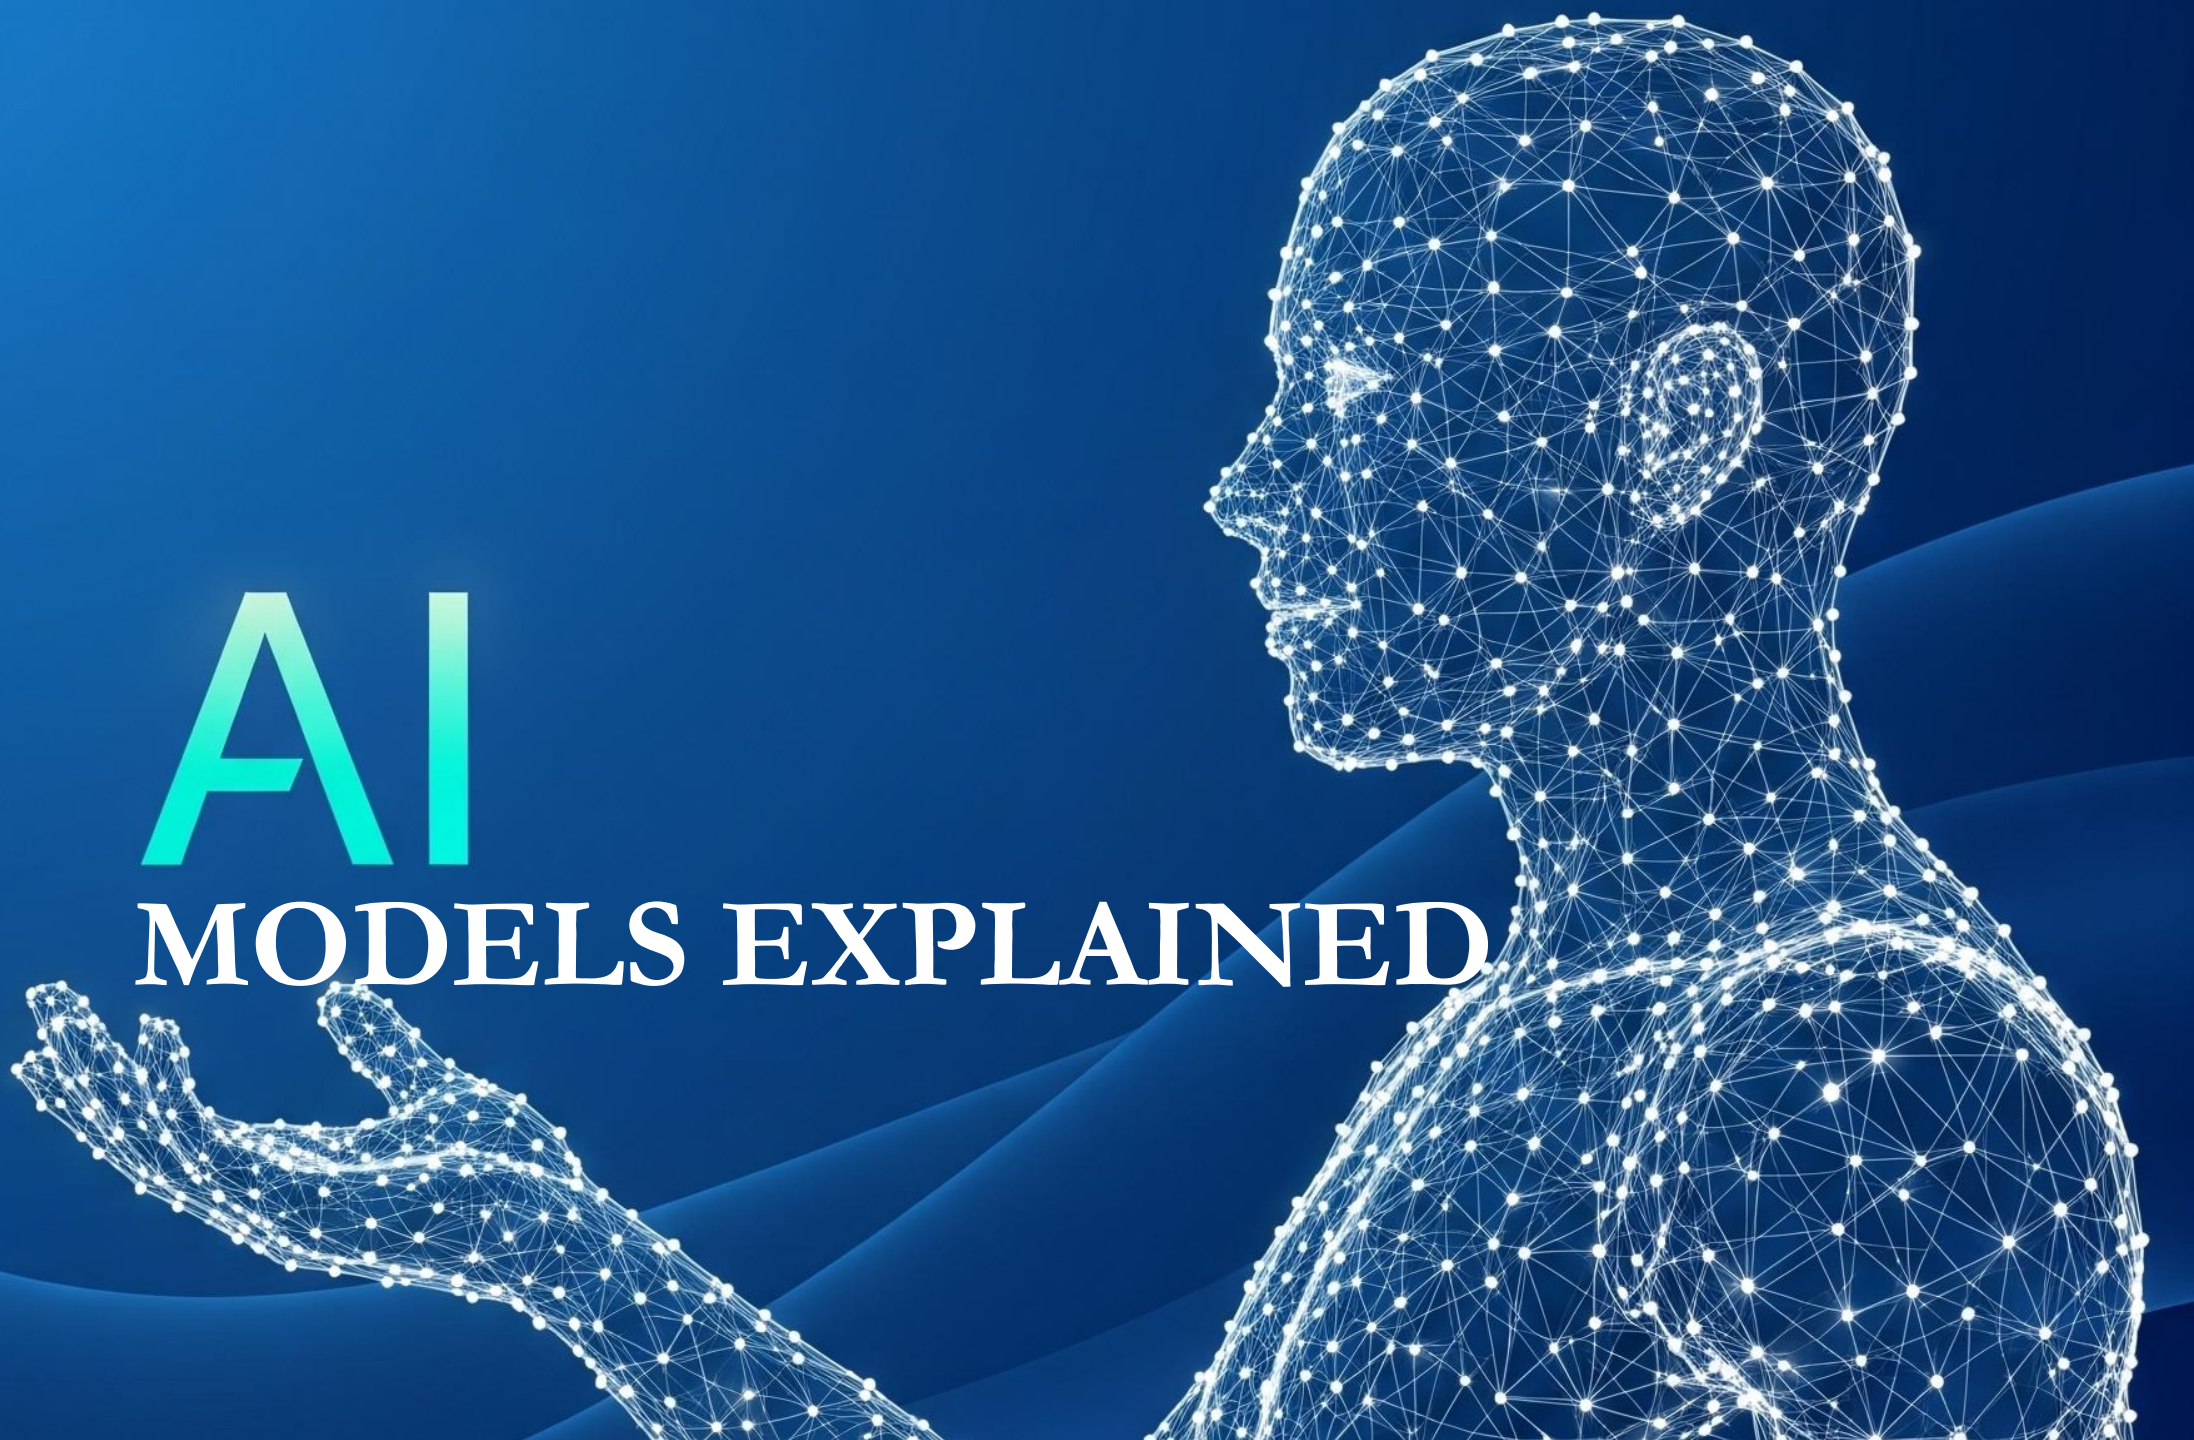

Artificial Intelligence (AI)

The "Umbrella" Term: This represents the entire field. It covers any computer program or algorithm that mimics human intelligence.

What it includes: Everything from a simple logic program to the most advanced sci-fi robots.

Key Takeaway: If a vendor sells you "AI," they haven't told you much yet. It's like saying they are selling you "food"—it could be an apple, or it could be a 5-course meal.

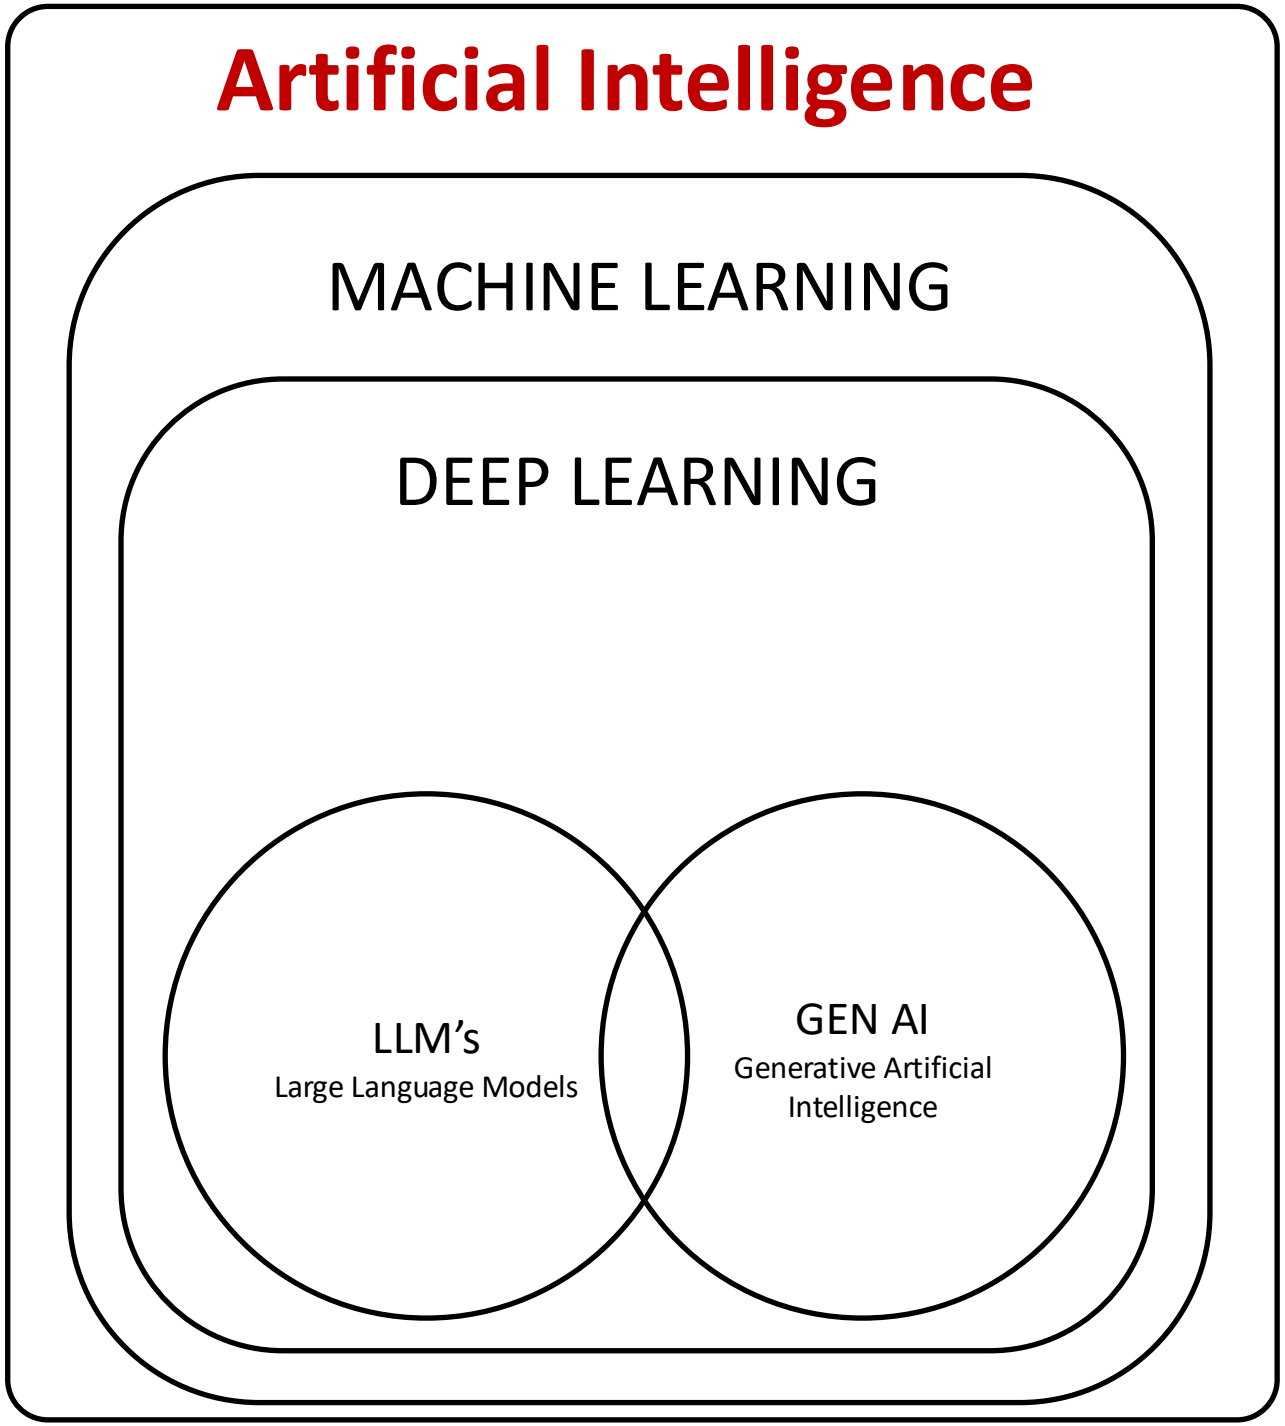

Artificial Intelligence (AI)

The Machine Learning (ML)

ML: This is a subset of AI. Instead of following strict rules programmed by a human, these systems learn from data.

How it works: You feed it Excel sheets of historical data, and it figures out the patterns on its own.

Gov Example: Predicting which roads will need maintenance next year based on traffic data.

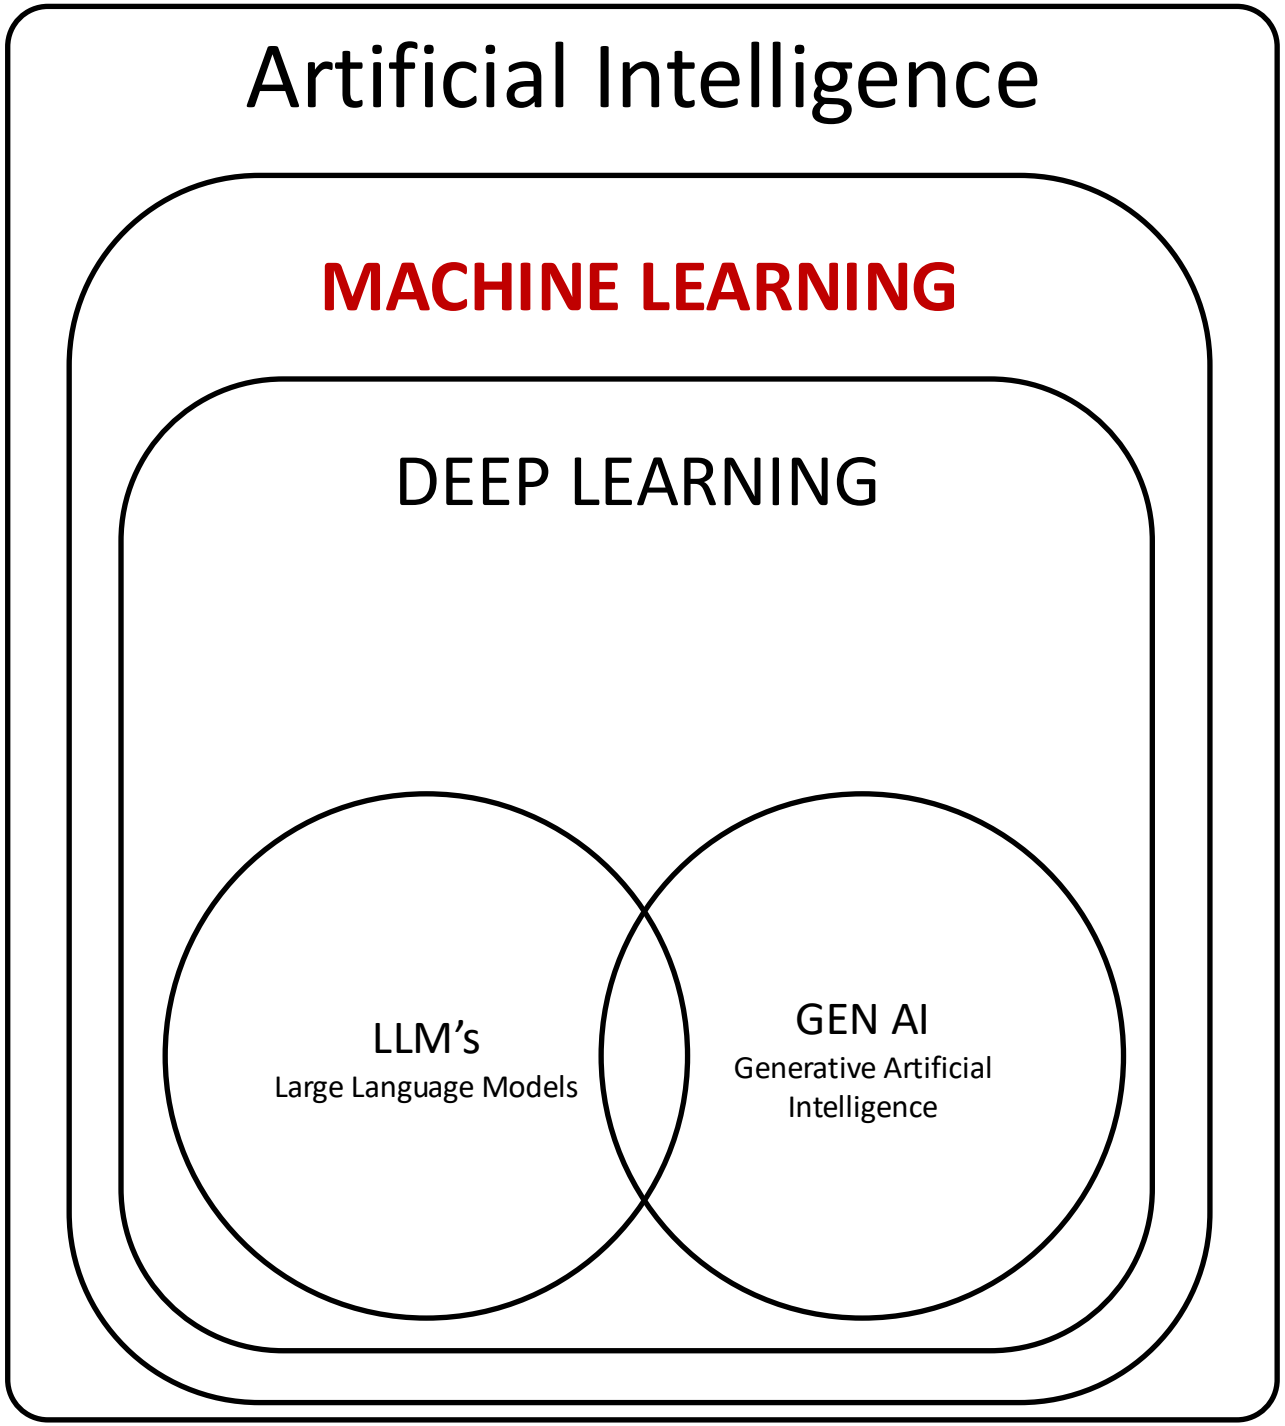

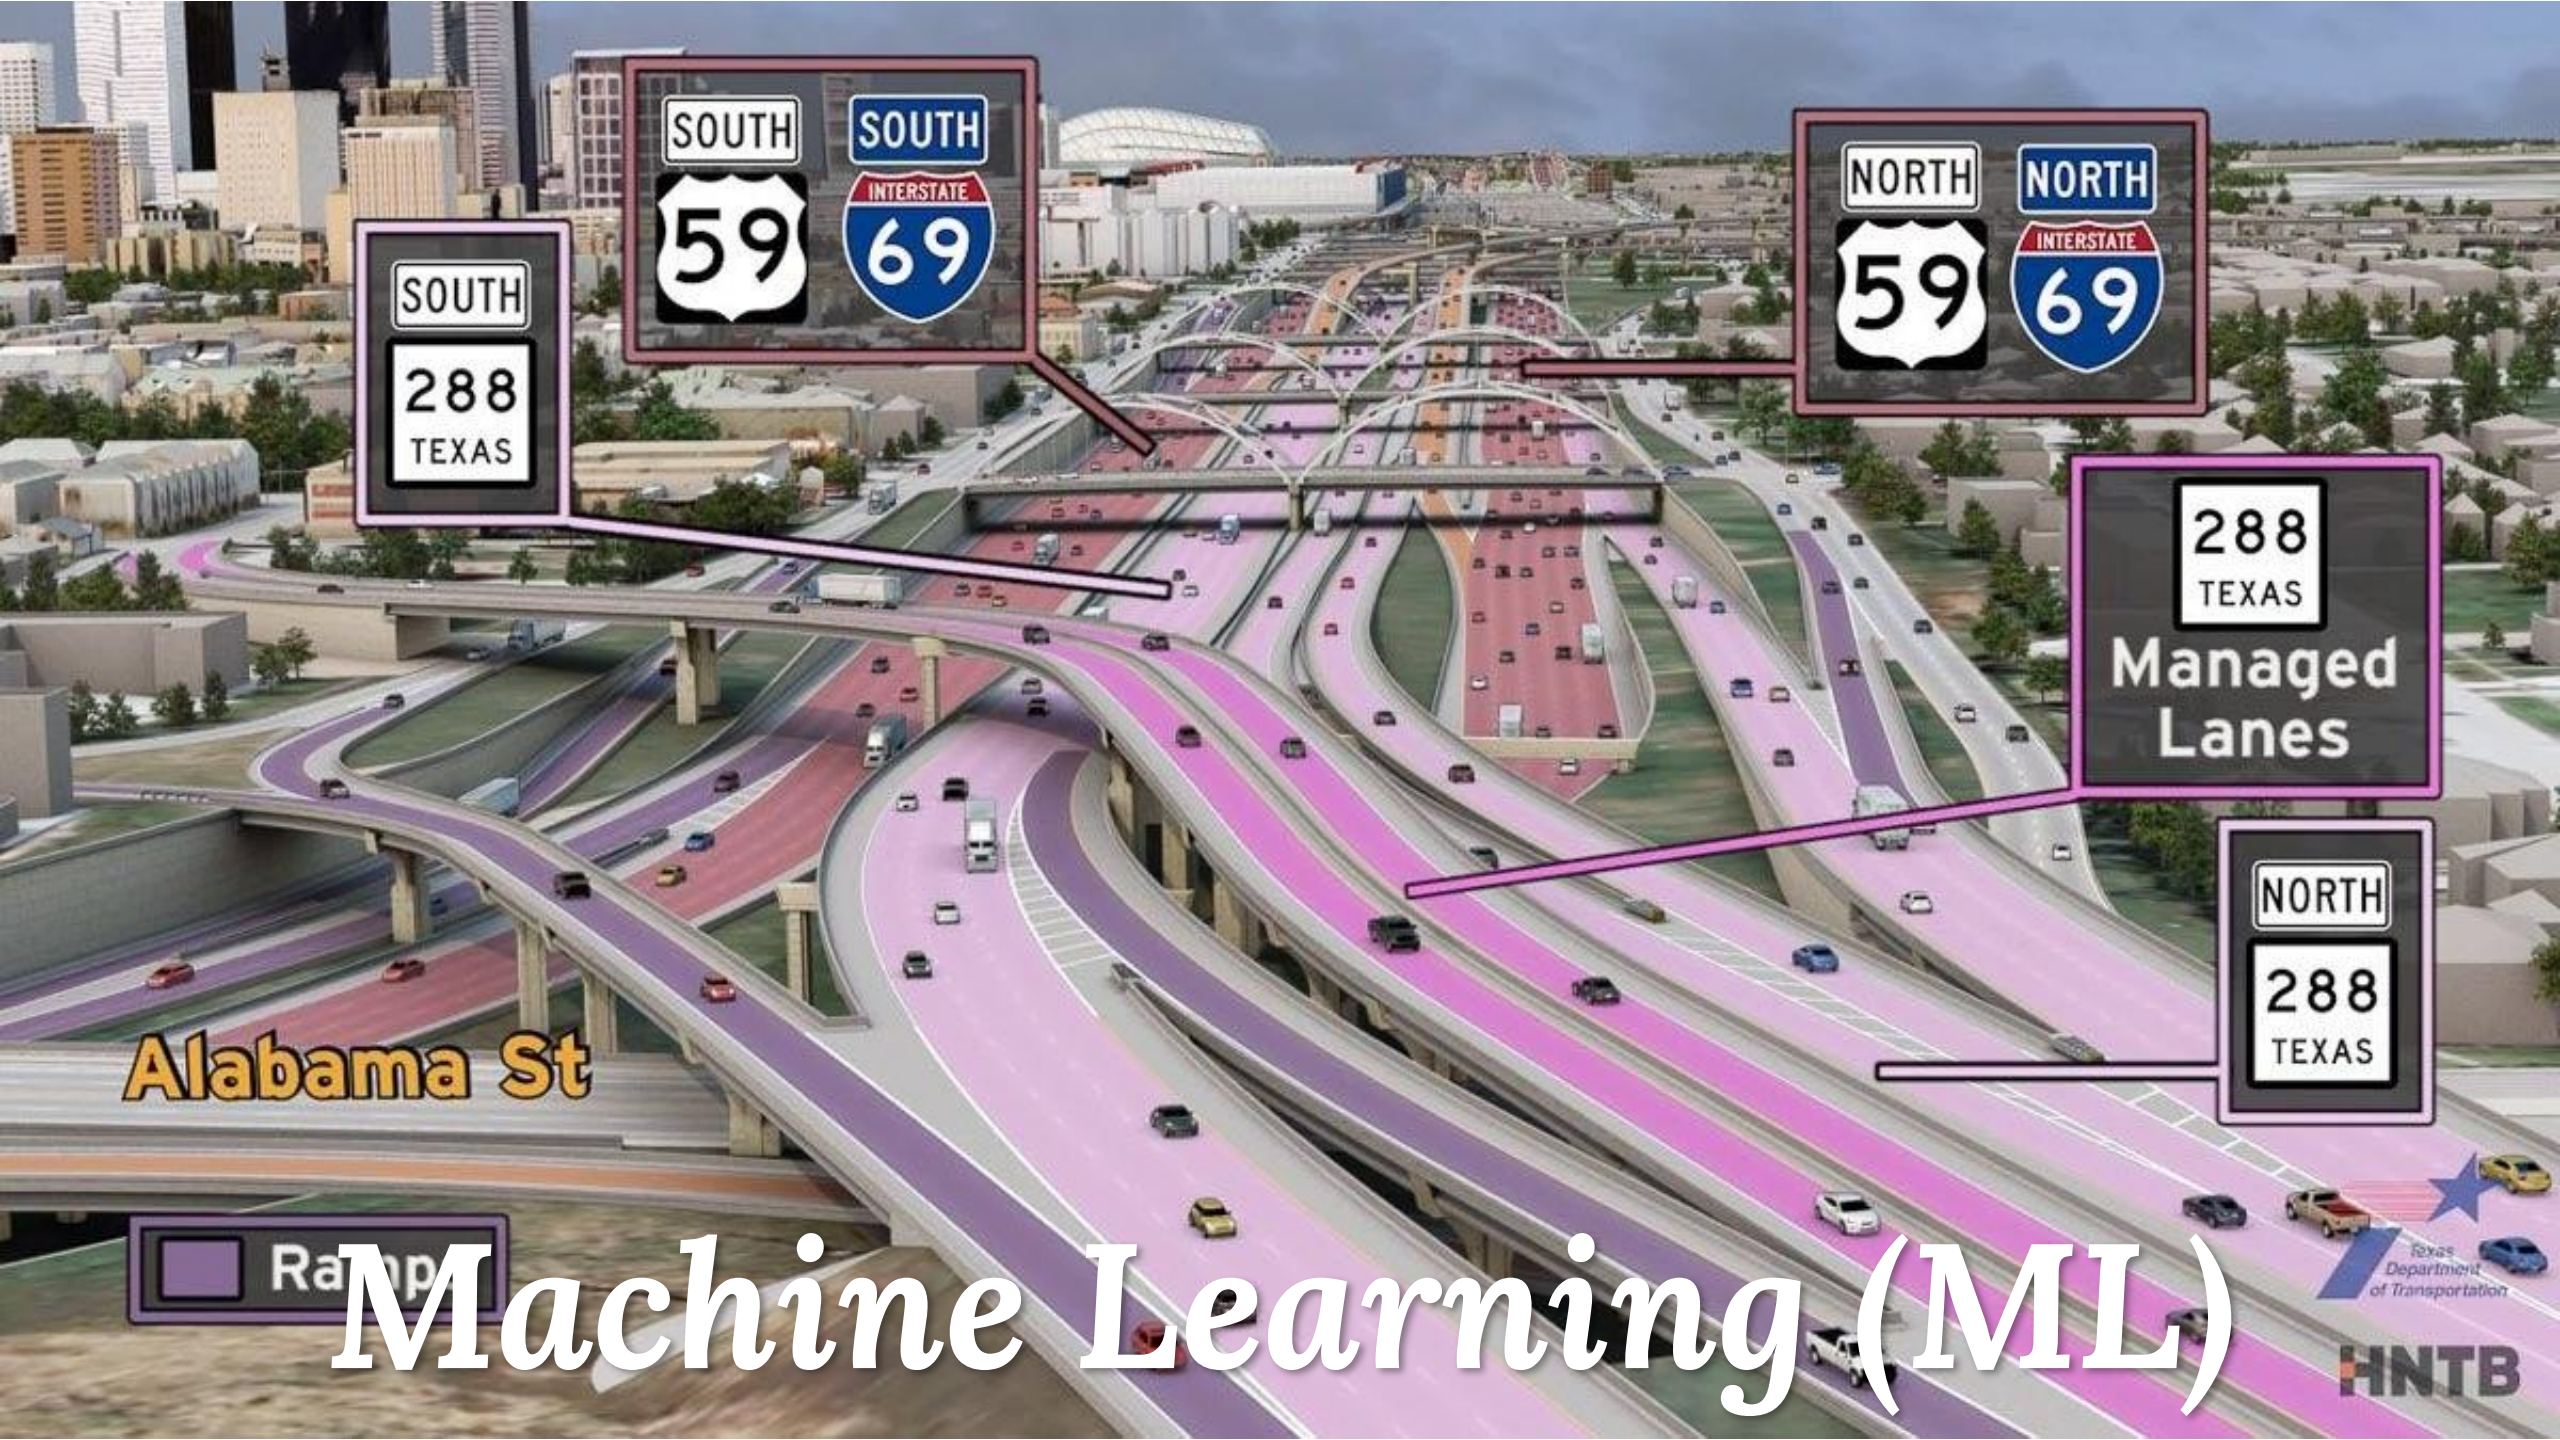

Machine Learning (ML)

Deep Learning (DL)

The "Brain": This is a specialized, more powerful subset of Machine Learning. It uses "***Neural Networks***"—layers of math designed to mimic the human brain's structure.

What makes it different: Standard Machine Learning is good with spreadsheets (numbers). Deep Learning is required when the data is messy and complex, like images, voice, and video.

Gov Example: Facial recognition at customs or analyzing drone footage for search and rescue.

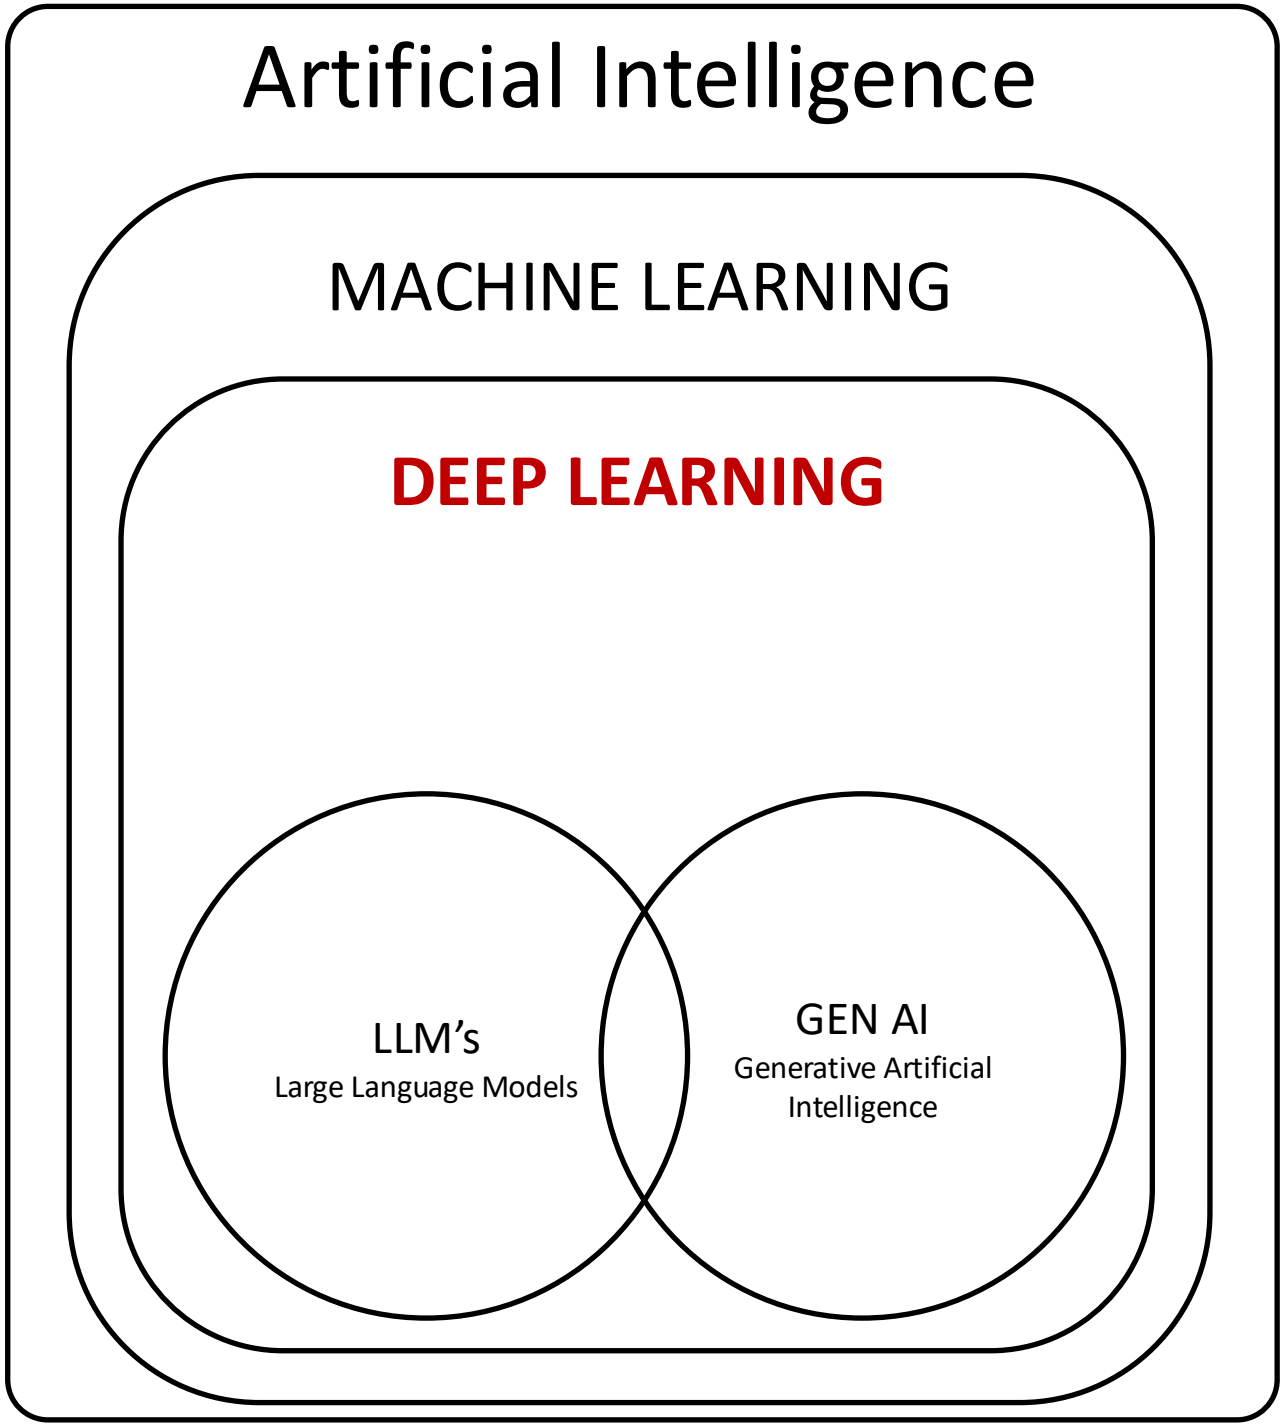

Deep Learning (DL)

GEN AI TOOLS

Gemini

OpenAI

Meta

Grok

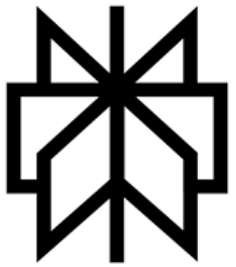

Perplexity

Apple Intelligence

Claude

Artificial Intelligence

MACHINE LEARNING

DEEP LEARNING

LLM's
Large Language Models

GEN AI
Generative Artificial Intelligence

AI-powered software designed to work independently with minimal setup.

STAND ALONE AI TOOLS

Gemini

OpenAI

Grok

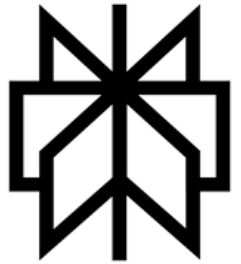

Perplexity

Claude

Established software platforms that have embedded Artificial Intelligence directly into their existing workflows

TOOLS WITH INTEGRATED AI FEATURES

Microsoft 365 Copilot

Gemini +

An AI application that is tailor-made to solve a specific problem

CUSTOM AI SOLUTIONS

JOHNS HOPKINS
SCHOOL of EDUCATION

+

CUSTOM LLM

=

CUSTOM AI SOLUTION

60% Prediction Accuracy

HOW DOES AI WORK? (3 easy steps…)

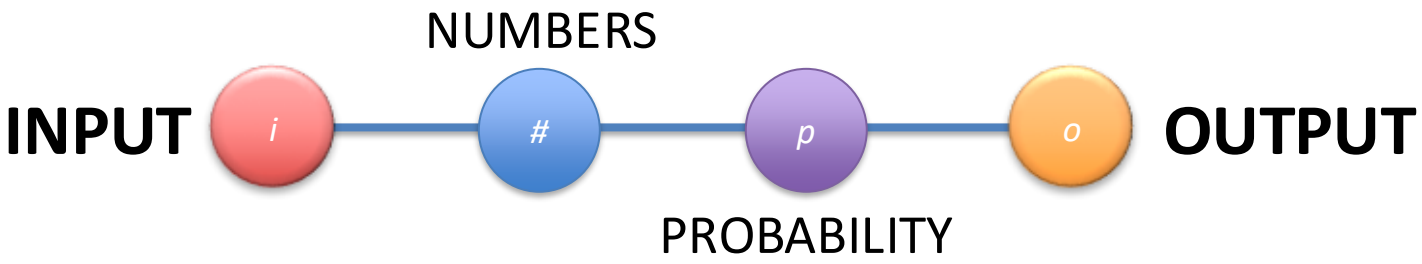

HOW DOES AI WORK? (2 easy steps…)

INPUT = What shape is the wheel?

#'s assigned = shape wheel

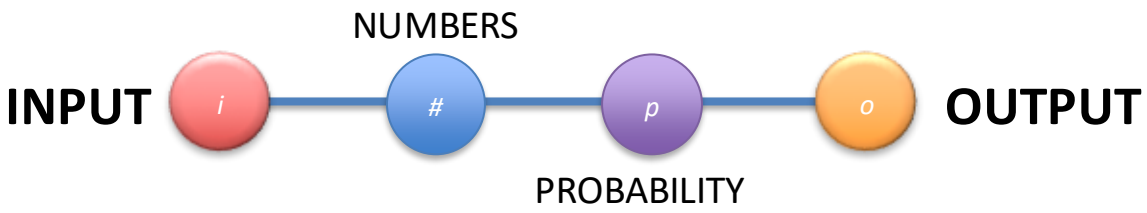

HOW DOES AI WORK? (1 easy step…)

Circle 97%

Rectangle 2%

Triangle 0.5%

Octagon 0.7%

probability determined

What shape is the wheel?

shape wheel

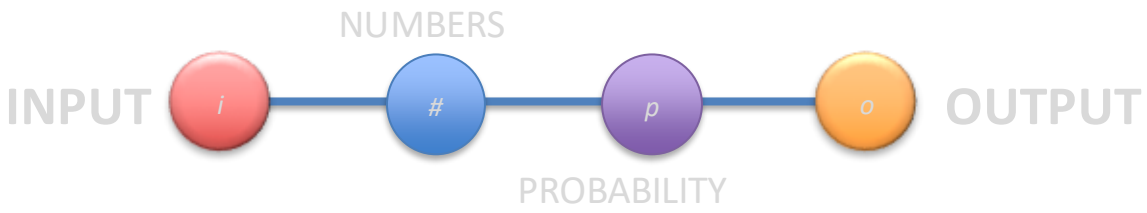

HOW DOES AI WORK?

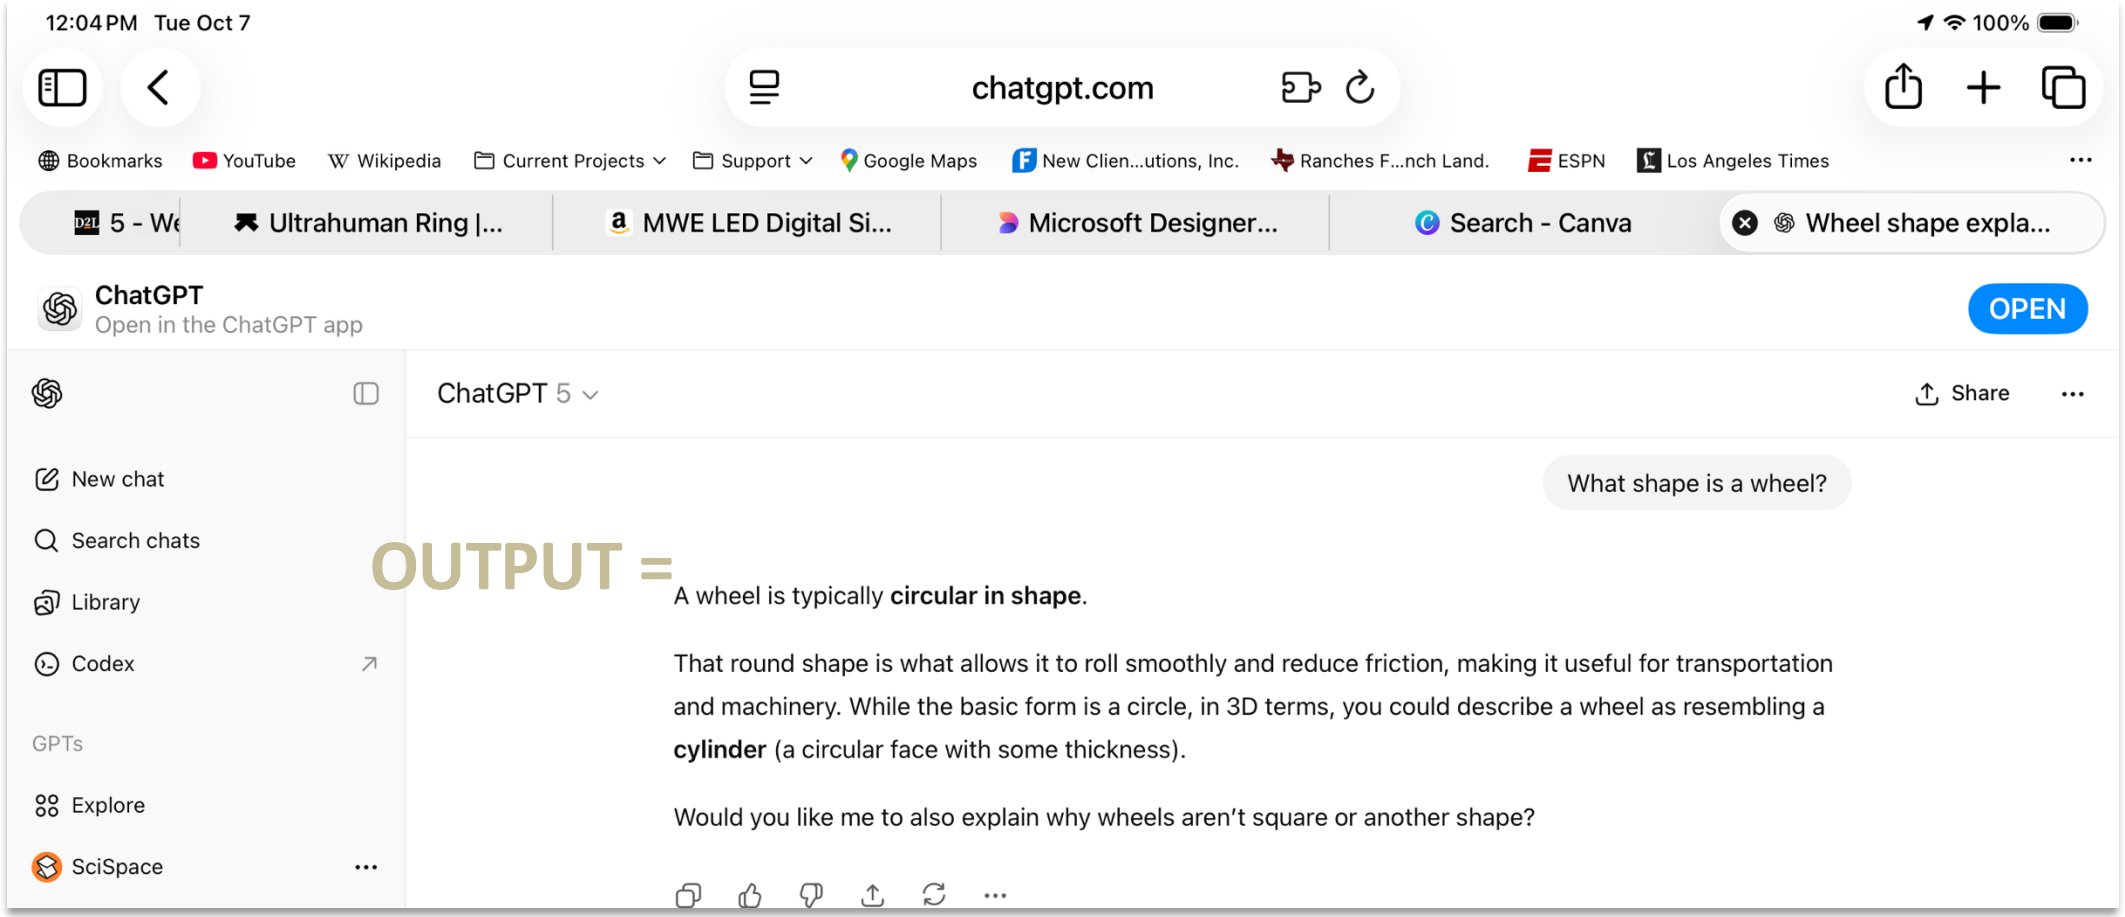

Artificial Intelligence

MACHINE LEARNING

SUPERVISED	UNSUPERVISED

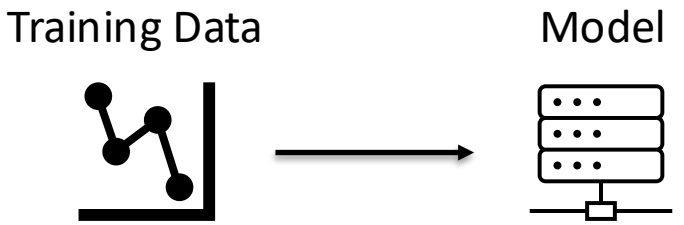

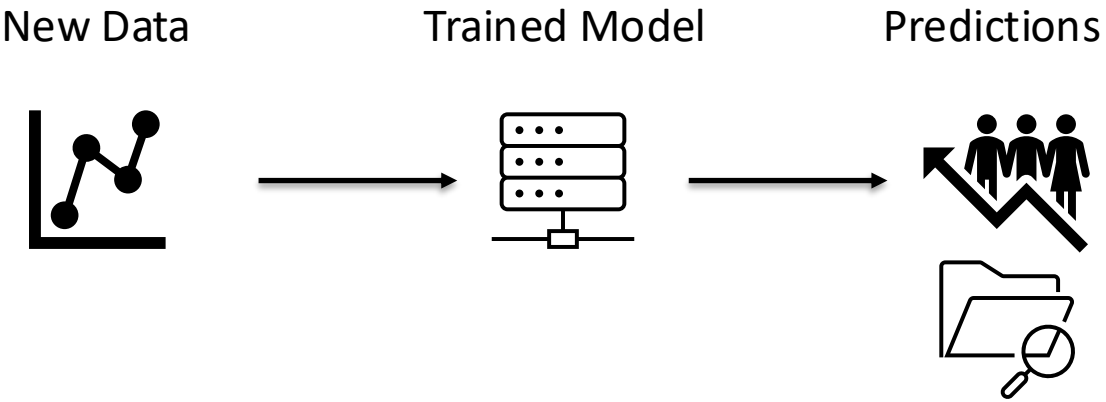

Supervised vs. Unsupervised ML

Feature	Supervised Learning	Unsupervised Learning
The Data	**Labeled** (Has an answer key)	**Unlabeled** (Raw data only)
The Analogy	A student taking a test with the teacher correcting them.	A librarian sorting books into genres without reading the titles.
Primary Goal	**Predict** future outcomes.	**Discover** hidden patterns.
Gov Use Case	Predicting which bridge needs maintenance.	Identifying new types of cyber-attacks.

Artificial Intelligence

MACHINE LEARNING

SUPERVISED	UNSUPERVISED

Nike Sales Data

Model

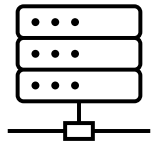

Adidas sales data

Trained Model

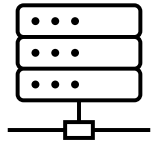

Predictions

Artificial Intelligence

MACHINE LEARNING

DEEP LEARNING

DISCRIMINITIVE

LLM's
Large Language Models

AI
Artificial Intelligence

+ Deep Learning Model = Fraud

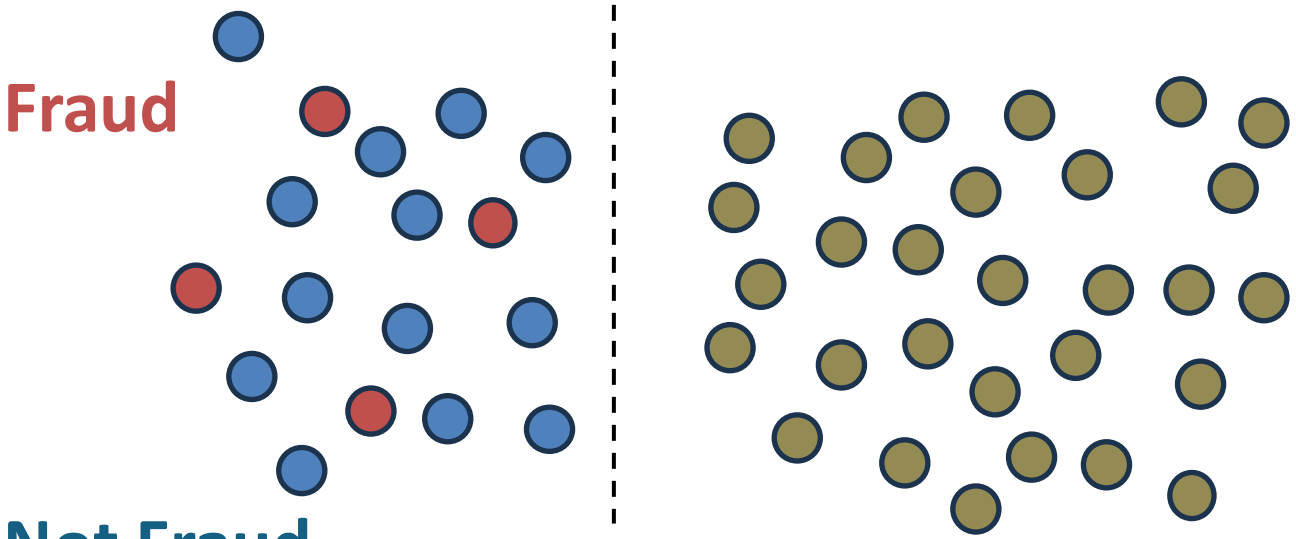

(labeled Data)

(unlabeled Data)

Algorithms: K-means clustering |K-nearest neighbor | Regression

+ Deep Learning Model = Fraud

MODEL MAKES PREDICTIONS

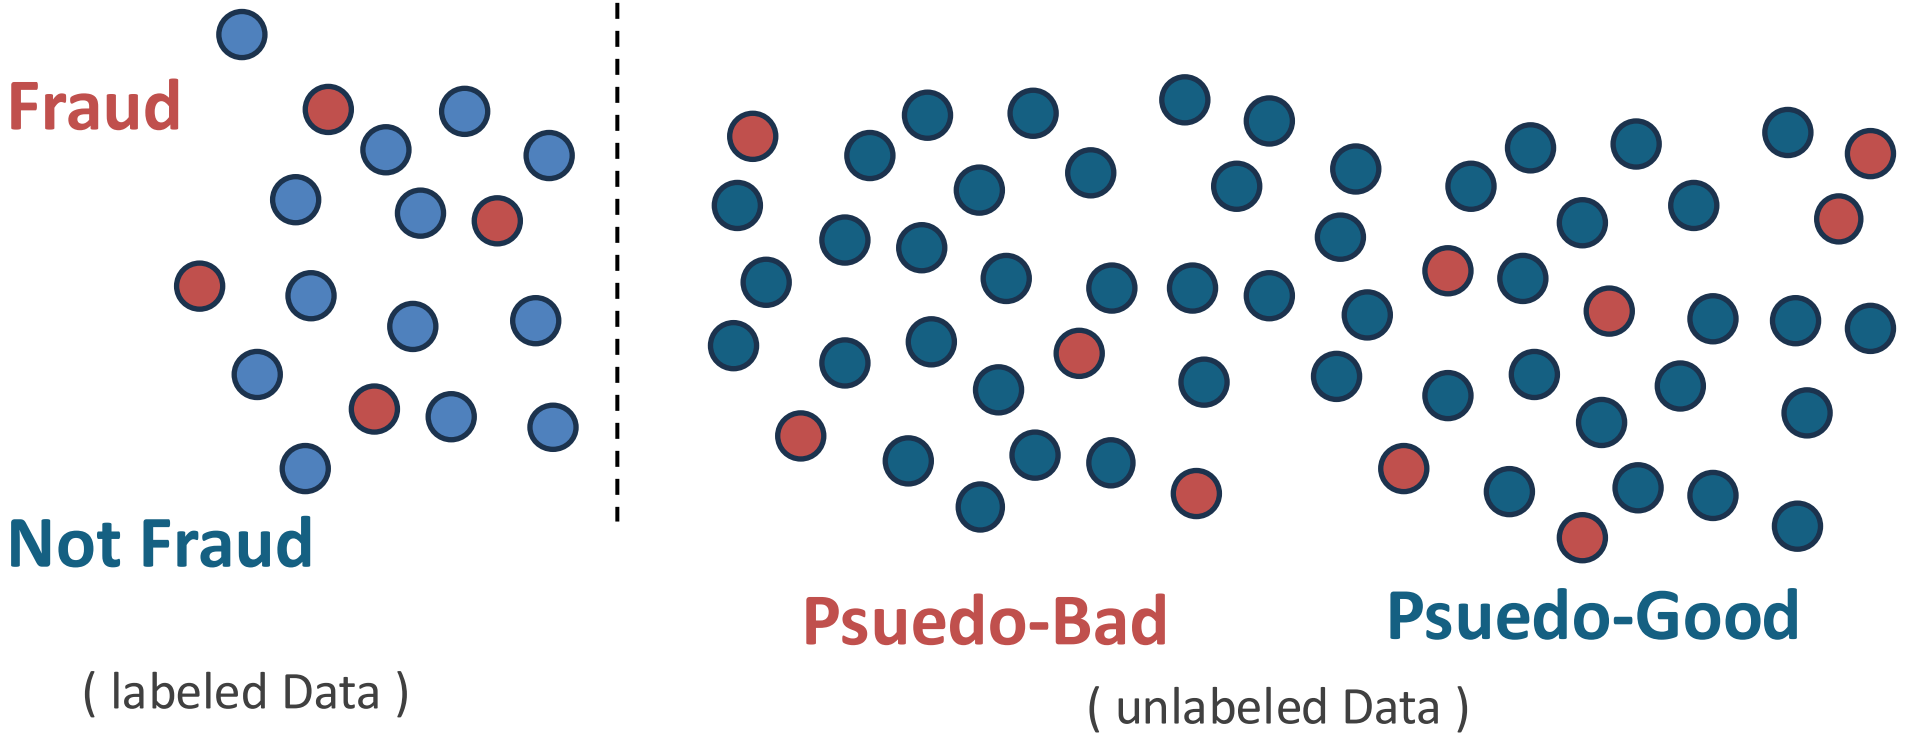

DISCRIMINITIVE MODEL EXAMPLE

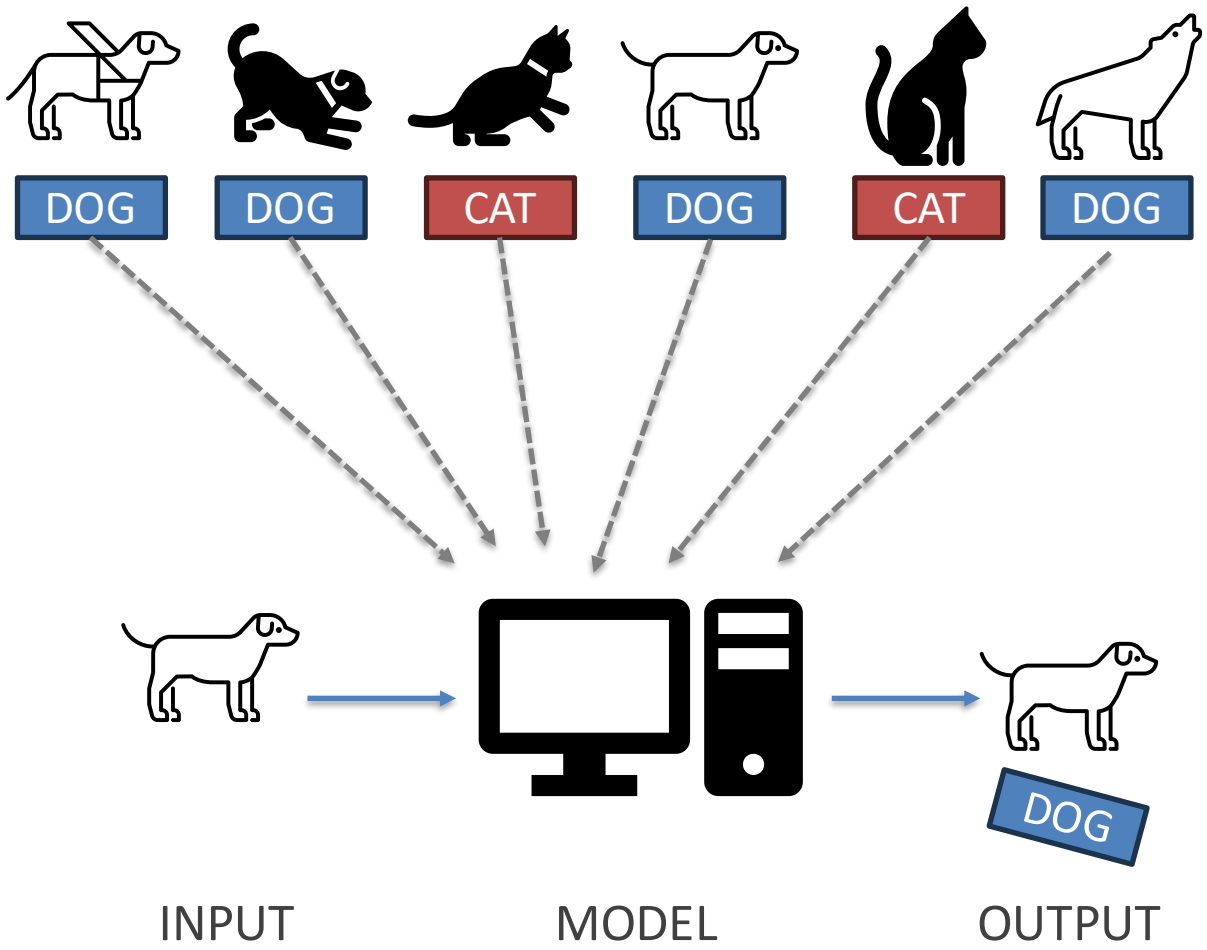

Summary for Leaders: When you ask AI to find a "Cat," it isn't looking for a pet. It is scanning millions of pixels for **Triangles (ears)**, **Vertical Slits (eyes)**, and **Short Snouts**. If the math adds up, it labels it "Cat."

Artificial Intelligence

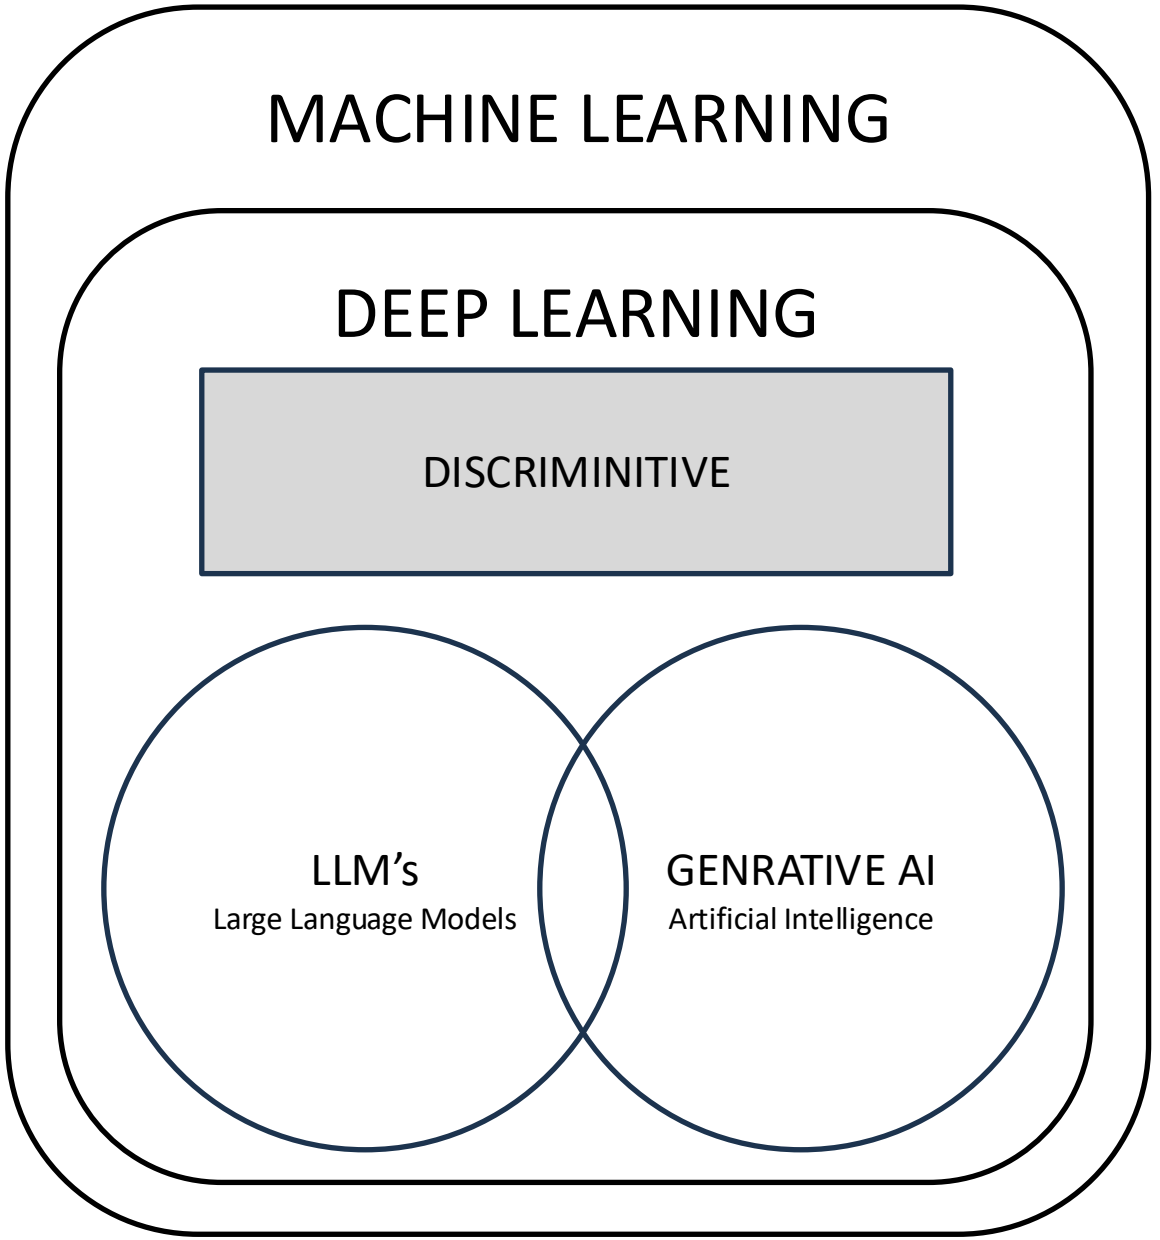

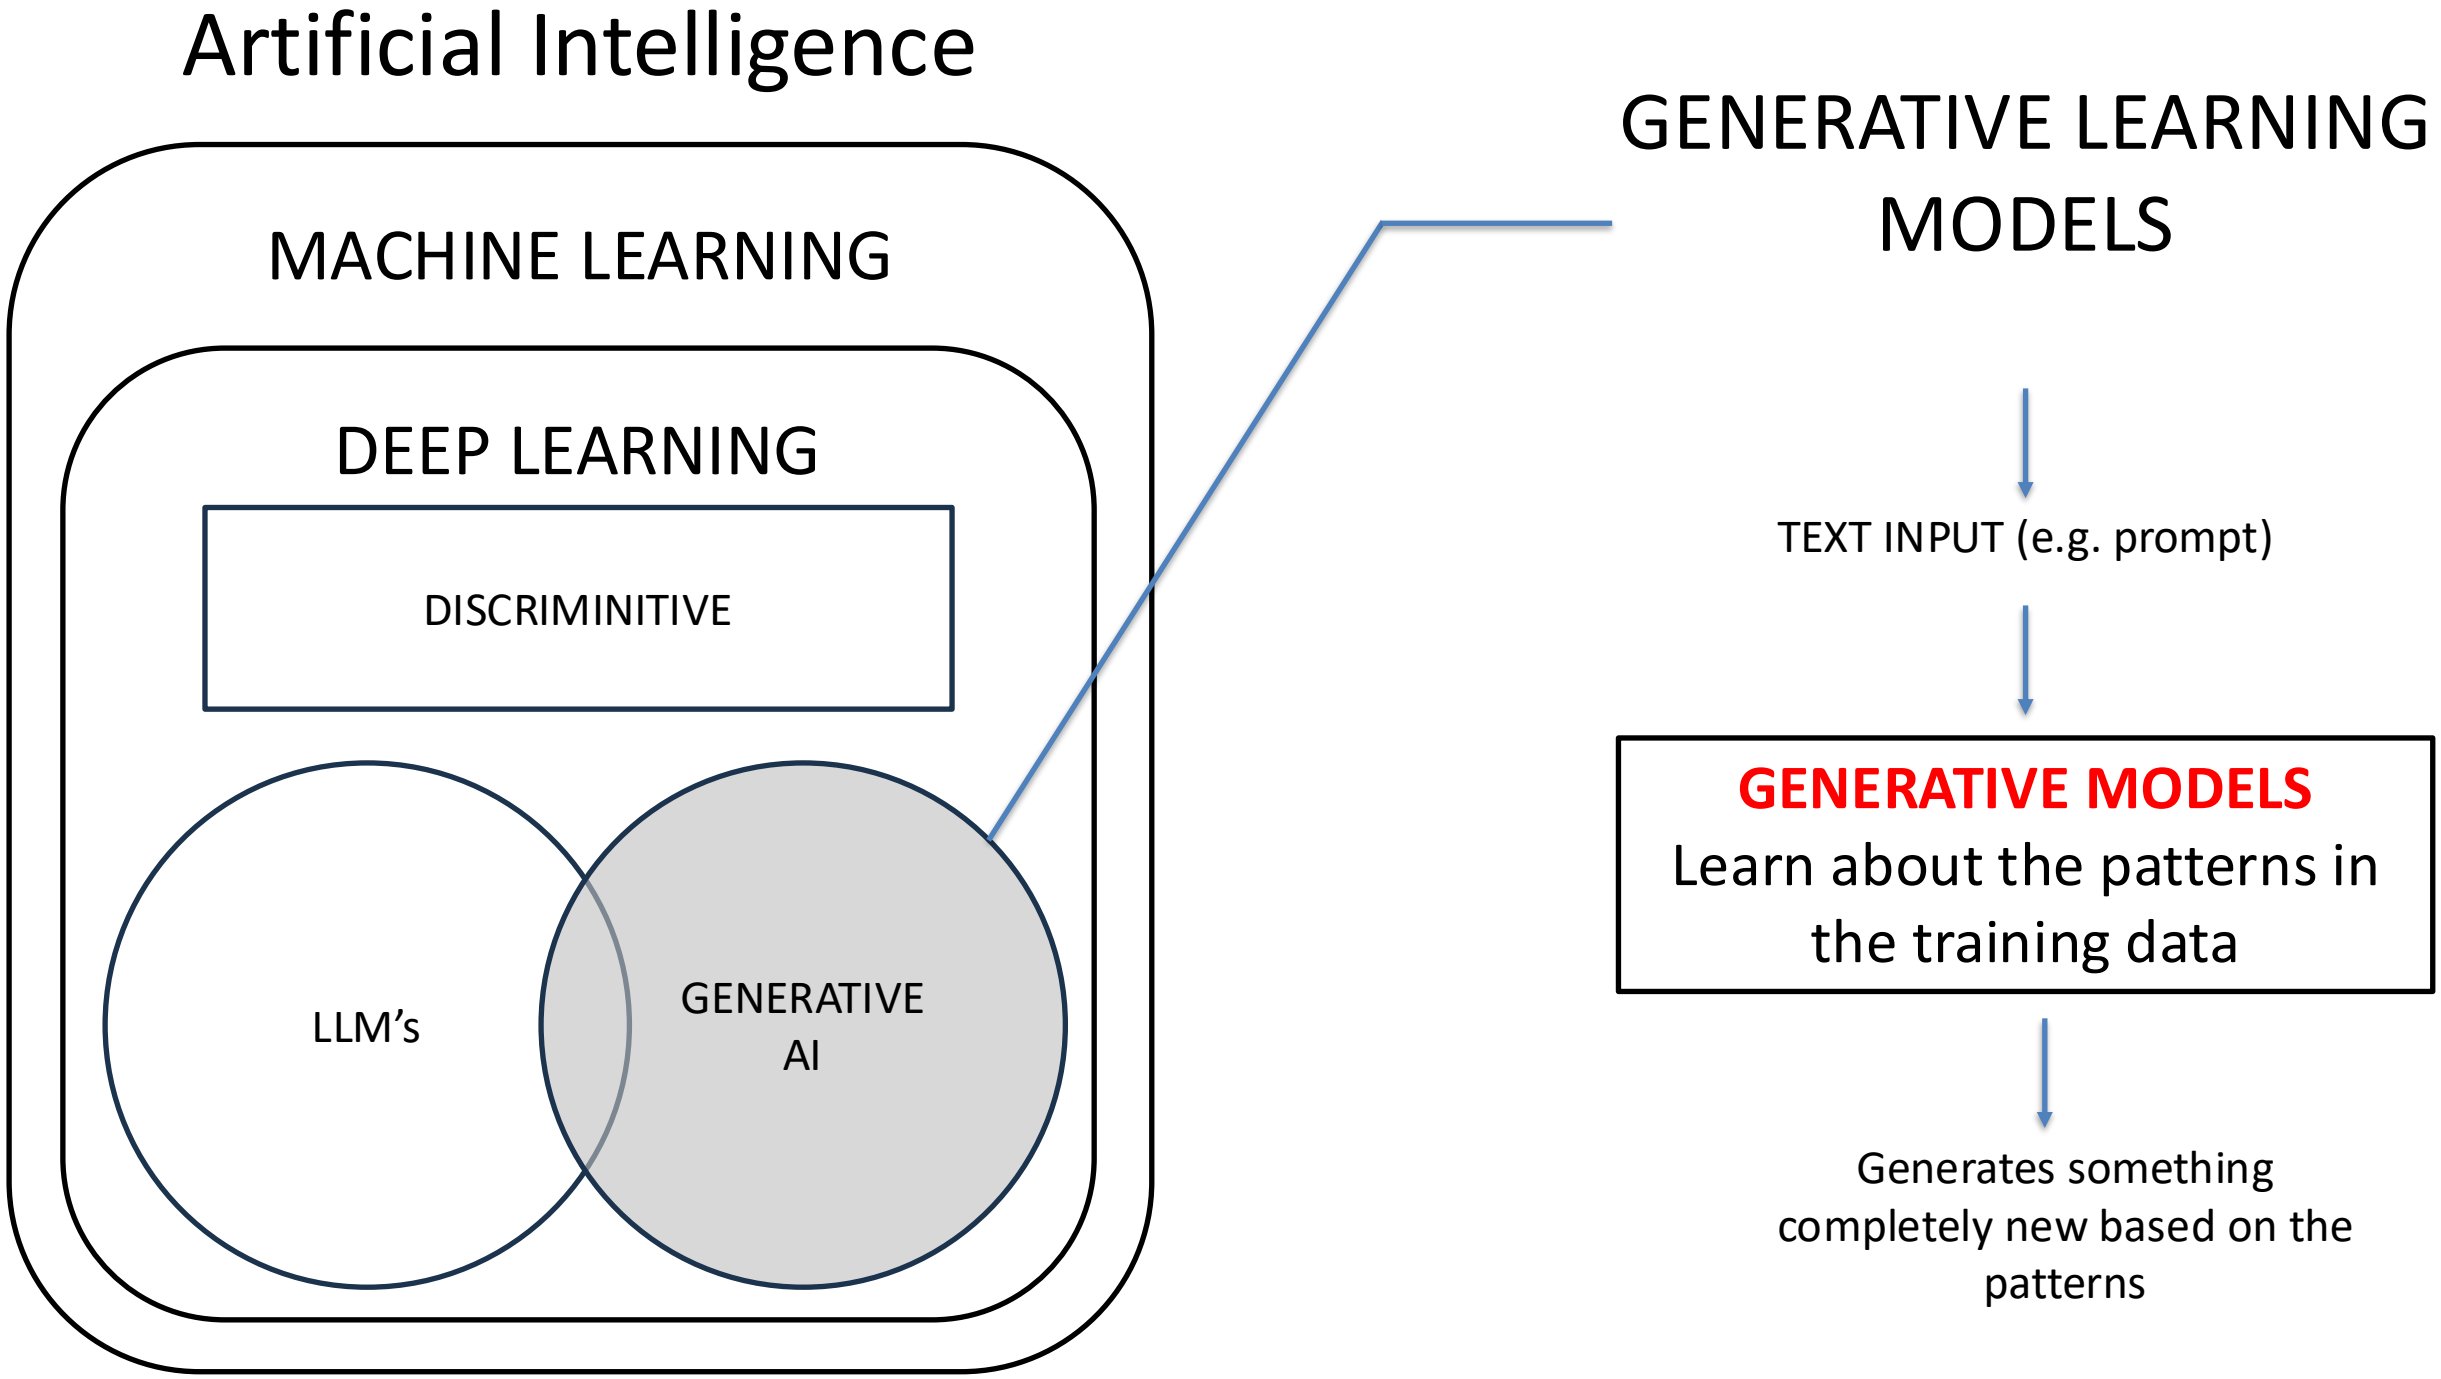
Artificial Intelligence
MACHINE LEARNING
DEEP LEARNING
DISCRIMINITIVE
LLM's
GENERATIVE AI
GENERATIVE LEARNING MODELS
TEXT INPUT (e.g. prompt)
GENERATIVE MODELS
Learn about the patterns in the training data
Generates something completely new based on the patterns

GENERATIVE MODEL EXAMPLE

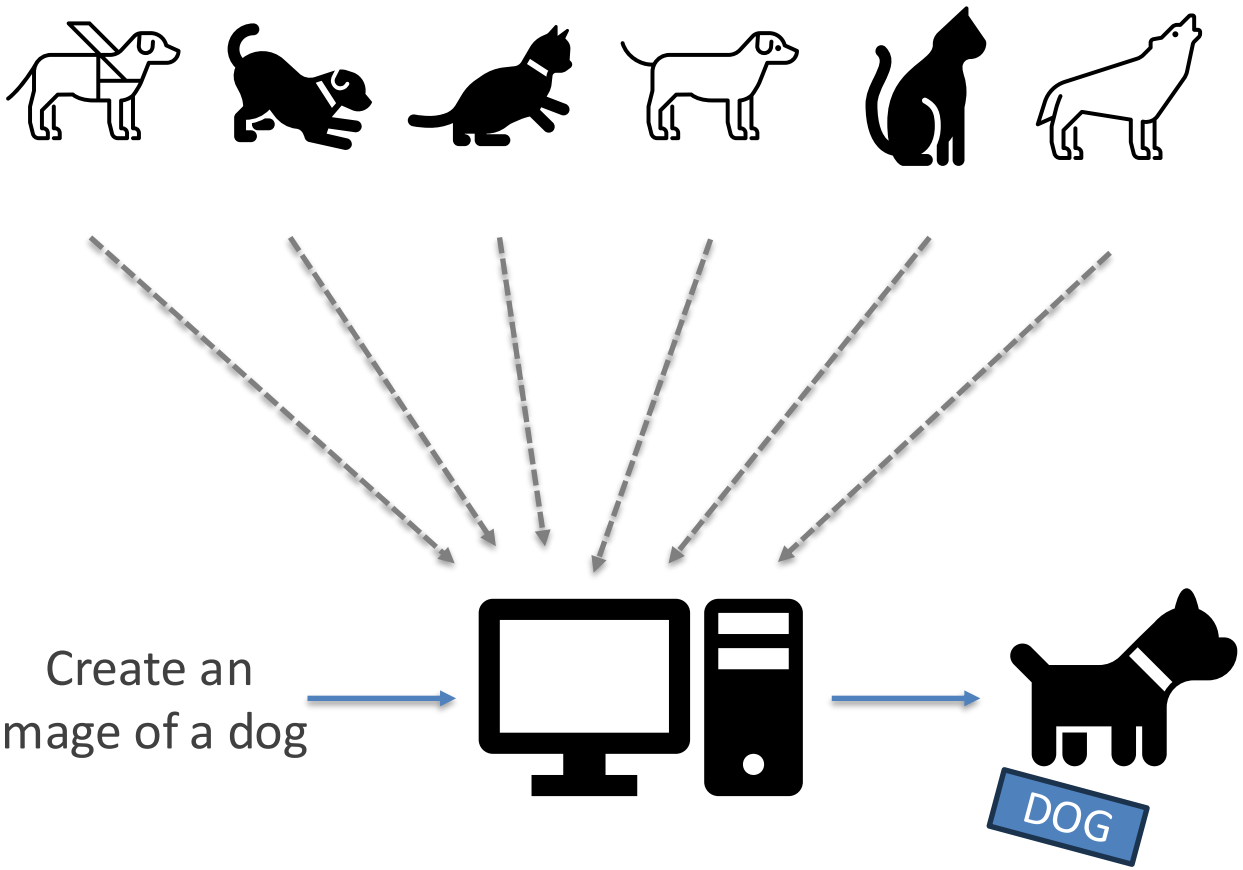

Artificial Intelligence

MACHINE LEARNING

DEEP LEARNING

DISCRIMINITIVE

LLM's
Large Language Models

GENERATIVE AI
Artificial Intelligence

The AI difference between a DOG & CAT

1. The "Triangle" Rule (Ears)

The Cat Pattern: Almost all domestic cats have rigid, triangular, upright ears. To an AI, finding two triangles on top of a circle (the head) is a **high-probability indicator** of a cat.

The Dog Pattern: Dogs are mathematically messy. Their ears can be floppy, cropped, long, short, or upright. If the AI detects "floppy" shapes on the head, it immediately increases the probability score for "Dog."

2. The "Snout" Vector (Face Geometry)

The Cat Pattern: Cats generally have flat faces. The distance between the eyes and the nose is short.

The Dog Pattern: Most dogs have a protruding snout. The AI measures the "depth" and length of the nose bridge. A longer nose bridge usually equals "Dog."

> *AI Confusion Point:* Pugs and French Bulldogs often confuse early AI models because their flat faces mimic the mathematical pattern of a cat!

3. The "Slit" vs. "Circle" (Eyes)

The Cat Pattern: Under bright light, cat pupils contract into vertical slits. This is a distinct vertical line feature that AI edge-detectors pick up easily.

The Dog Pattern: Dogs have round pupils.

If the image resolution is high enough, the AI scans the eyes specifically for that vertical line feature.

GENERATIVE AI MODEL TYPES

Text-to-text	Take natural language input and produce text output
Text-to-image	Trained on large set of images, and generate new images
Text-to-video	Can generate and edit videos
Text-to-3D	Can be used to produce gaming assets
Text-to-Task	Trained to perform a specific task or action based on text input

ChatGPT

Google Bard

Google Bard

RULES

PATTERNS

MATCHING

Large Language Models (LLM)

Think of a Large Language Model

(the engine behind ChatGPT)

as "**Autocorrect on Steroids**" combined with a "**Digital Librarian.**"

Large Language Models

It Predicts: Just like your phone guesses the next word when you text, an LLM is trained to predict the next word in a sentence. But instead of looking at just your last three words, it looks at billions of pages of text (books, laws, websites) to understand context.

It Generates: Because it has "read" so much, it doesn't just copy paste answers. It constructs new sentences word-by-word based on patterns it learned. If you ask for a speech in the style of MLK, it uses the mathematical patterns of MLK's speech to write something new, not just quote old speeches.

Large Language Models

The Problem: The Department of Veterans Affairs (VA) receives hundreds of thousands of feedback comments from veterans annually—handwritten notes, emails, and survey responses. Reading, tagging, and routing these manually takes thousands of staff hours, meaning critical issues (like a safety hazard at a clinic) might sit in a queue for weeks.

The LLM Solution: The VA implemented an LLM to "read" this incoming feedback.

- **What the LLM does:** It instantly categorizes the feedback (e.g., "Medical Billing," "Facility Cleanliness," "Homelessness Risk") and drafts a summary of the veteran's sentiment.
- **The Result:** Instead of staff spending 8 hours reading, they spend 1 hour reviewing the AI's tags. Critical feedback about suicide risk or immediate safety hazards is flagged instantly and routed to a human counselor, drastically reducing response time.

Large Language Models

You aren't replacing the caseworker; you are removing the "***paperwork bottleneck***" so the caseworker can spend more time actually helping veterans.

Types of AI (Patterns & Data)

Predictive AI

Identifies patterns to better predict future behavior.

(small amounts of data)

- Netflix
- YouTube
- Amazon

Generative AI

Take massive amounts of data and builds something valuable

(large amounts of data)

- Understand preferences
- Generate new images
- Generate Videos

Types of AI (Patterns & Data)

Predictive AI

(small amounts of data)

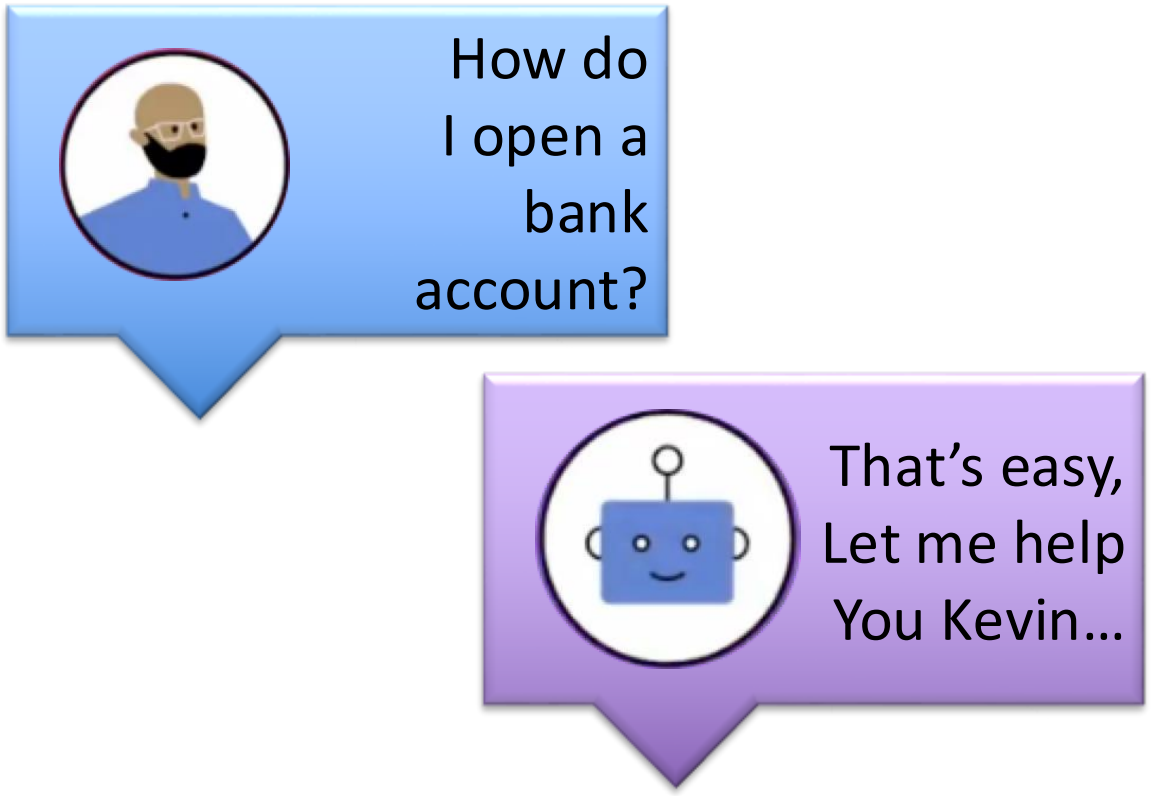

Generative AI

(large amounts of data)

Prompt: "Will today's financial news affect the balance on my accounts?"

Prompt: "Is this transaction tax deductible?"

Predictive Artificial Intelligence

NETFLIX

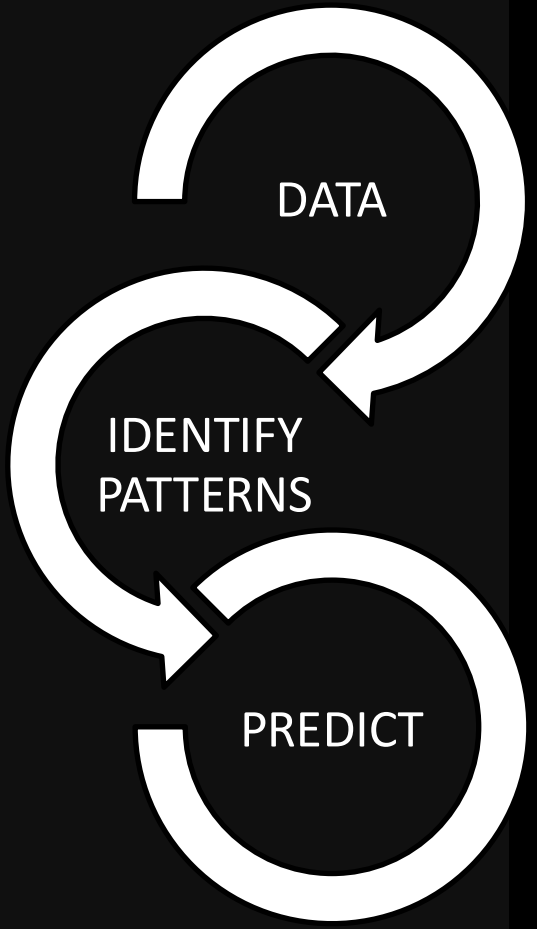

AI

Classification

Capability

3 - levels

Functionality

4 - types

AI Capability - Type 1

Artificial Narrow Intelligence (ANI) – Current State

Definition: AI that is excellent at one specific task but cannot do anything else. It mimics intelligence but has no consciousness.

Example: A chess computer can beat a grandmaster but cannot make a toast. Your agency's fraud detection algorithm is ANI—it knows math, but it doesn't know what "fraud" means.

AI Capability - Type 2

Artificial General Intelligence (AGI) – Theoretical

Definition: AI that equals human intellect. It can learn, solve problems, and plan for the future across any domain, not just the one it was trained on.

Example: An AI that could serve as a relentless, universally capable Chief of Staff—handling policy, HR, and strategy simultaneously with human level nuance.

Status: **Does not exist.** Estimates for its arrival range from 5 to 50 years.

AI Capability - Type 3

Artificial Super Intelligence (ASI) – Hypothetical

Definition: AI that vastly surpasses human intelligence in every field, from creativity to emotional intelligence.

Status: Pure science fiction at this moment.

AI Functionality (1 of 2)

How AI processes (data & memory)

Type 1: **Reactive Machines**

How it works: It has no memory. It looks at the world right now and makes a decision. It cannot learn from the past.

Example: The Deep Blue chess computer. It didn't "remember" previous games; it just calculated the best move for the current board layout.

Type 2: **Limited Memory – Current State**

How it works: It can look back at a short period of recent data to make decisions.

Example: Self-driving cars (they watch other cars' speed over time) and Chatbots/LLMs (they remember what you said 2 minutes ago in the chat). Note: Large Language Models like ChatGPT fall here. They appear smart, but their "memory" is limited to the conversation window.

AI Functionality (2 of 2)

How AI processes (data & memory)

Type 3: ***Theory of Mind – Future Goal***

How it works: The AI understands that humans have thoughts, emotions, and beliefs that affect their behavior. It can "empathize."

Example: A robotic nurse that recognizes a patient is frightened (not just elevated heart rate) and changes its tone of voice to comfort them.

Type 4: **Self-Awareness – The End Goal**

How it works: The AI knows it exists. It has consciousness, desires, and feelings.

Example: The robots in movies like I, Robot or Ex Machina.

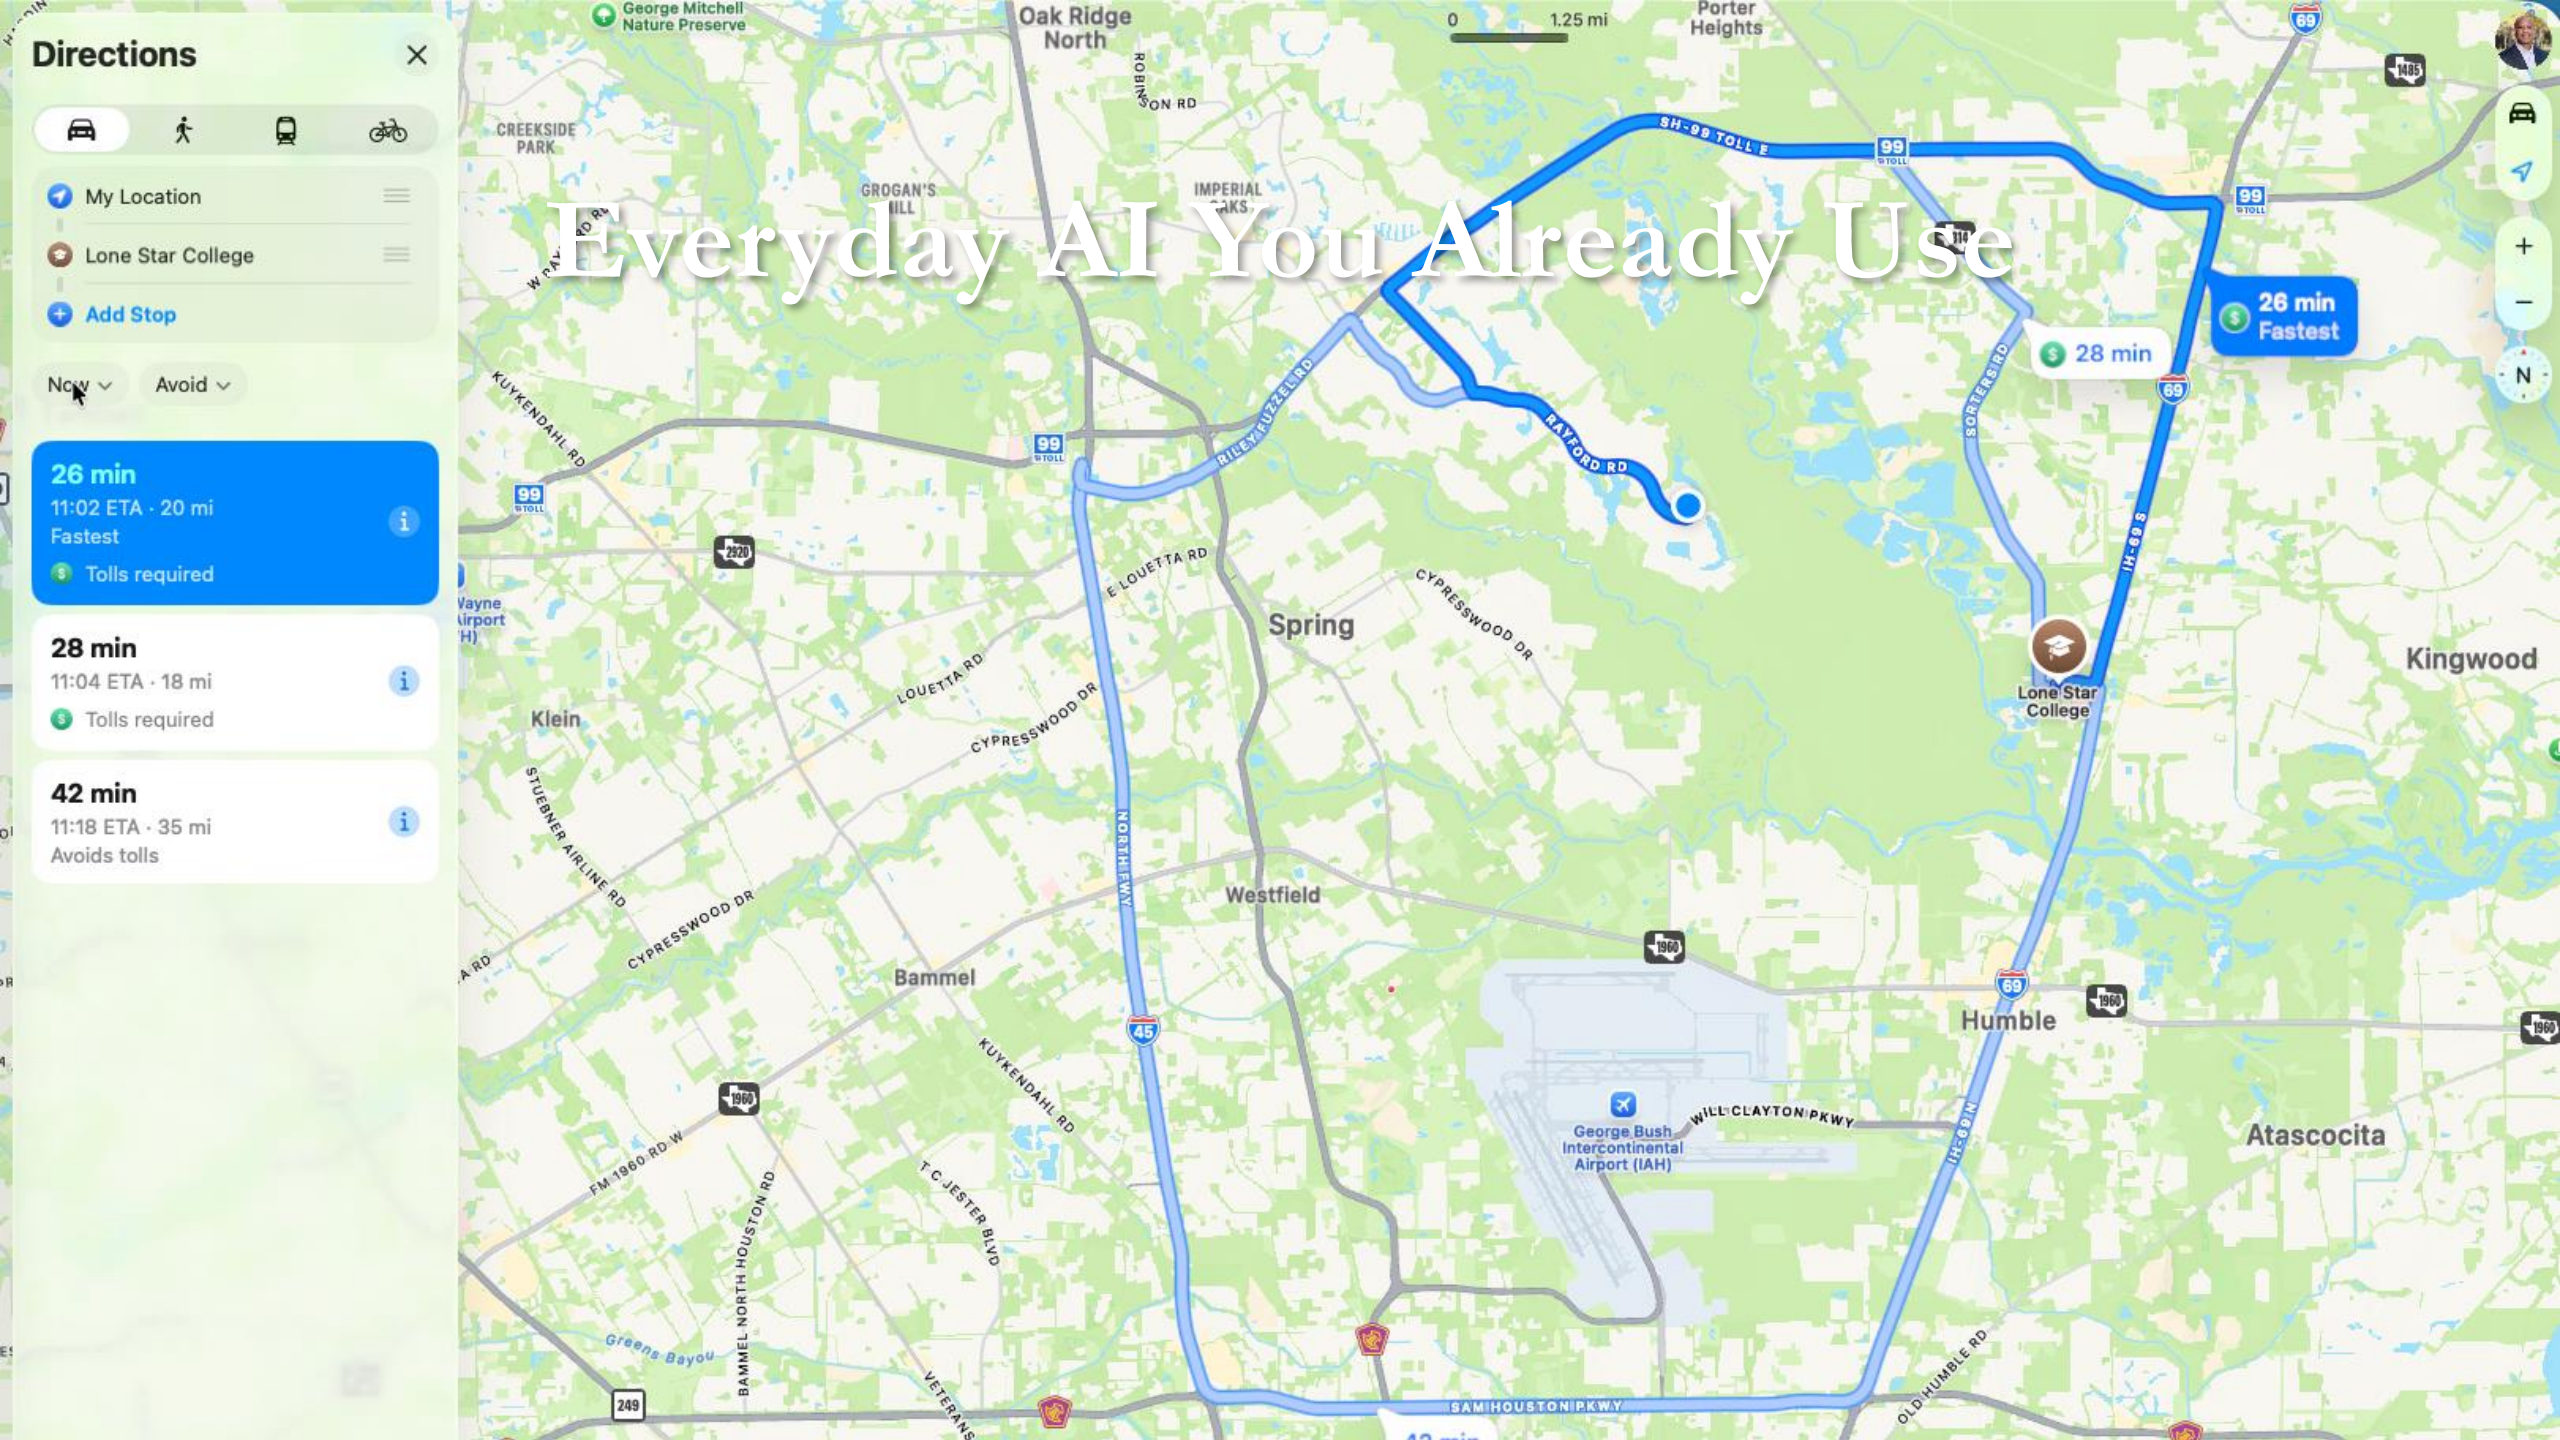
Directions
My Location
Lone Star College
Add Stop
Now
Avoid
26 min
11:02 ETA · 20 mi
Fastest
Tolls required
28 min
11:04 ETA · 18 mi
Tolls required
42 min
11:18 ETA · 35 mi
Avoids tolls
Everyday AI You Already Use
26 min
Fastest
28 min
Lone Star College
Spring
Klein
Bammel
Westfield
Humble
Kingwood
Atascocita
George Bush Intercontinental Airport (IAH)
Oak Ridge North
Porter Heights
SH-99 TOLL E
RILEY FUZZEL RD
RAYFORD RD
SORTERS RD
IH-69 S
IH-69 N
E LOUETTA RD
LOUETTA RD
CYPRESSWOOD DR
NORTH FWY
KUYKENDAHL RD
STUEBNER AIRLINE RD
FM 1960 RD W
BAMMEL NORTH HOUSTON RD
T C JESTER BLVD
WILL CLAYTON PKWY
SAM HOUSTON PKWY
OLD HUMBLE RD
Greens Bayou
George Mitchell Nature Preserve
CREEKSIDE PARK
GROGAN'S MILL
IMPERIAL OAKS
0
1.25 mi
N

Everyday AI You Already Use

Everyday AI You Already Use

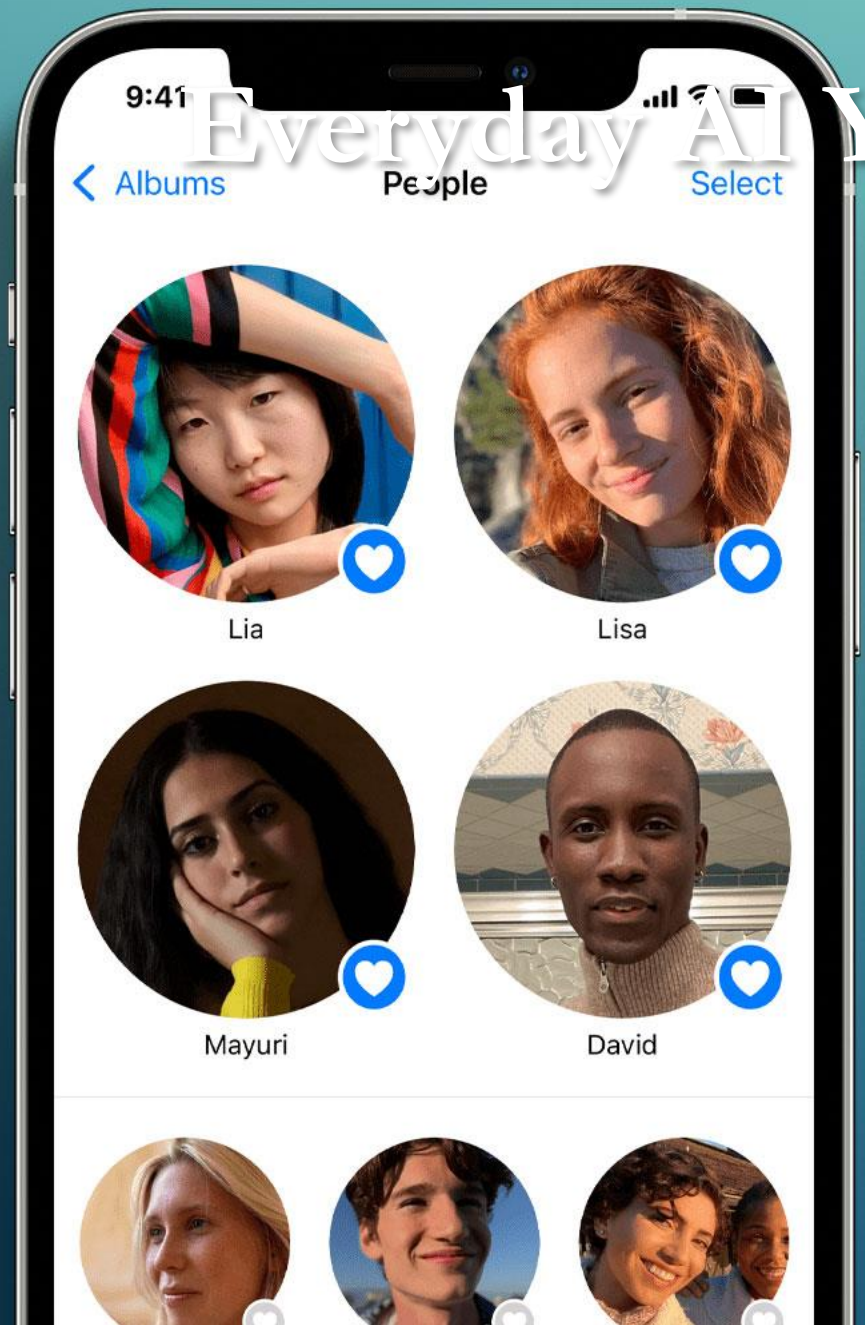

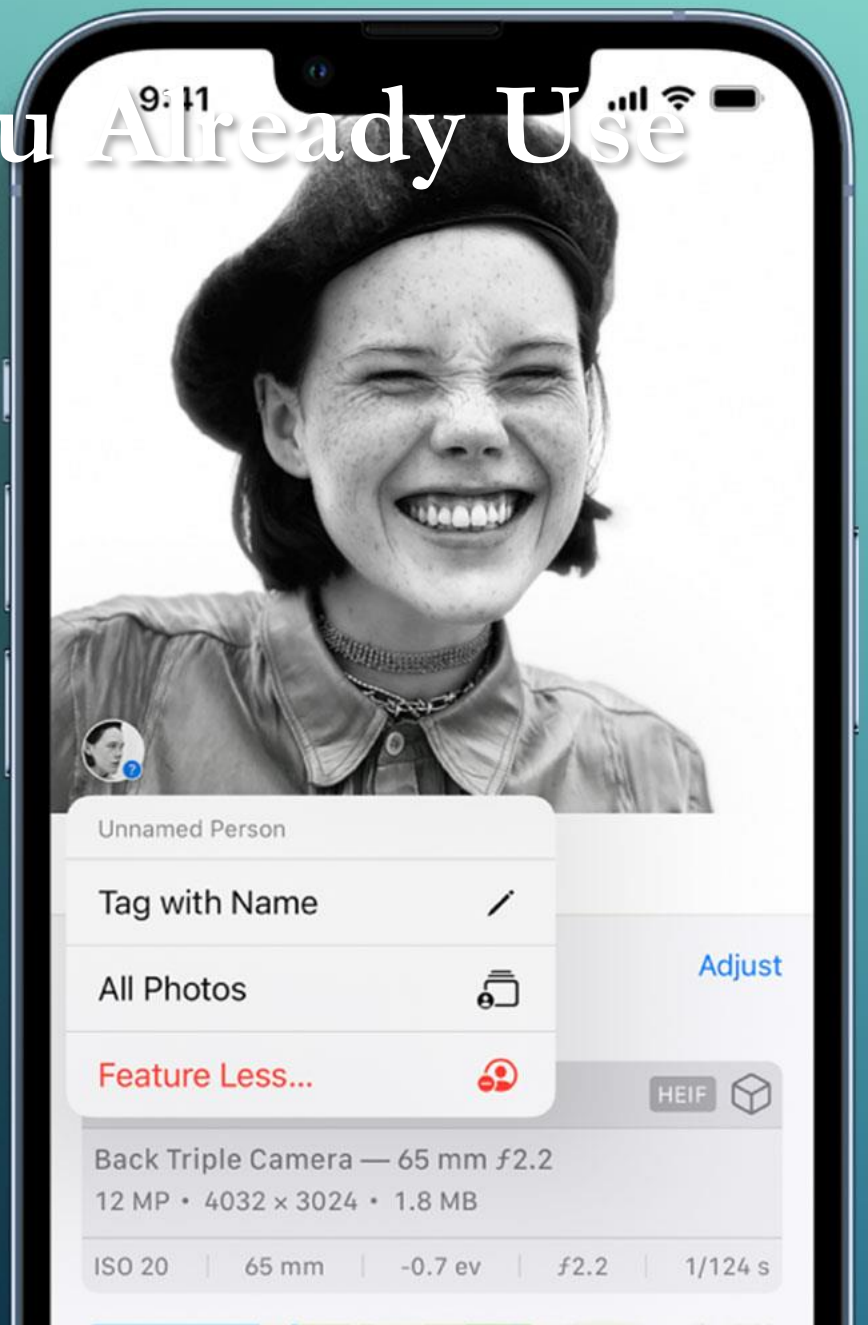

Everyday AI You Already Use

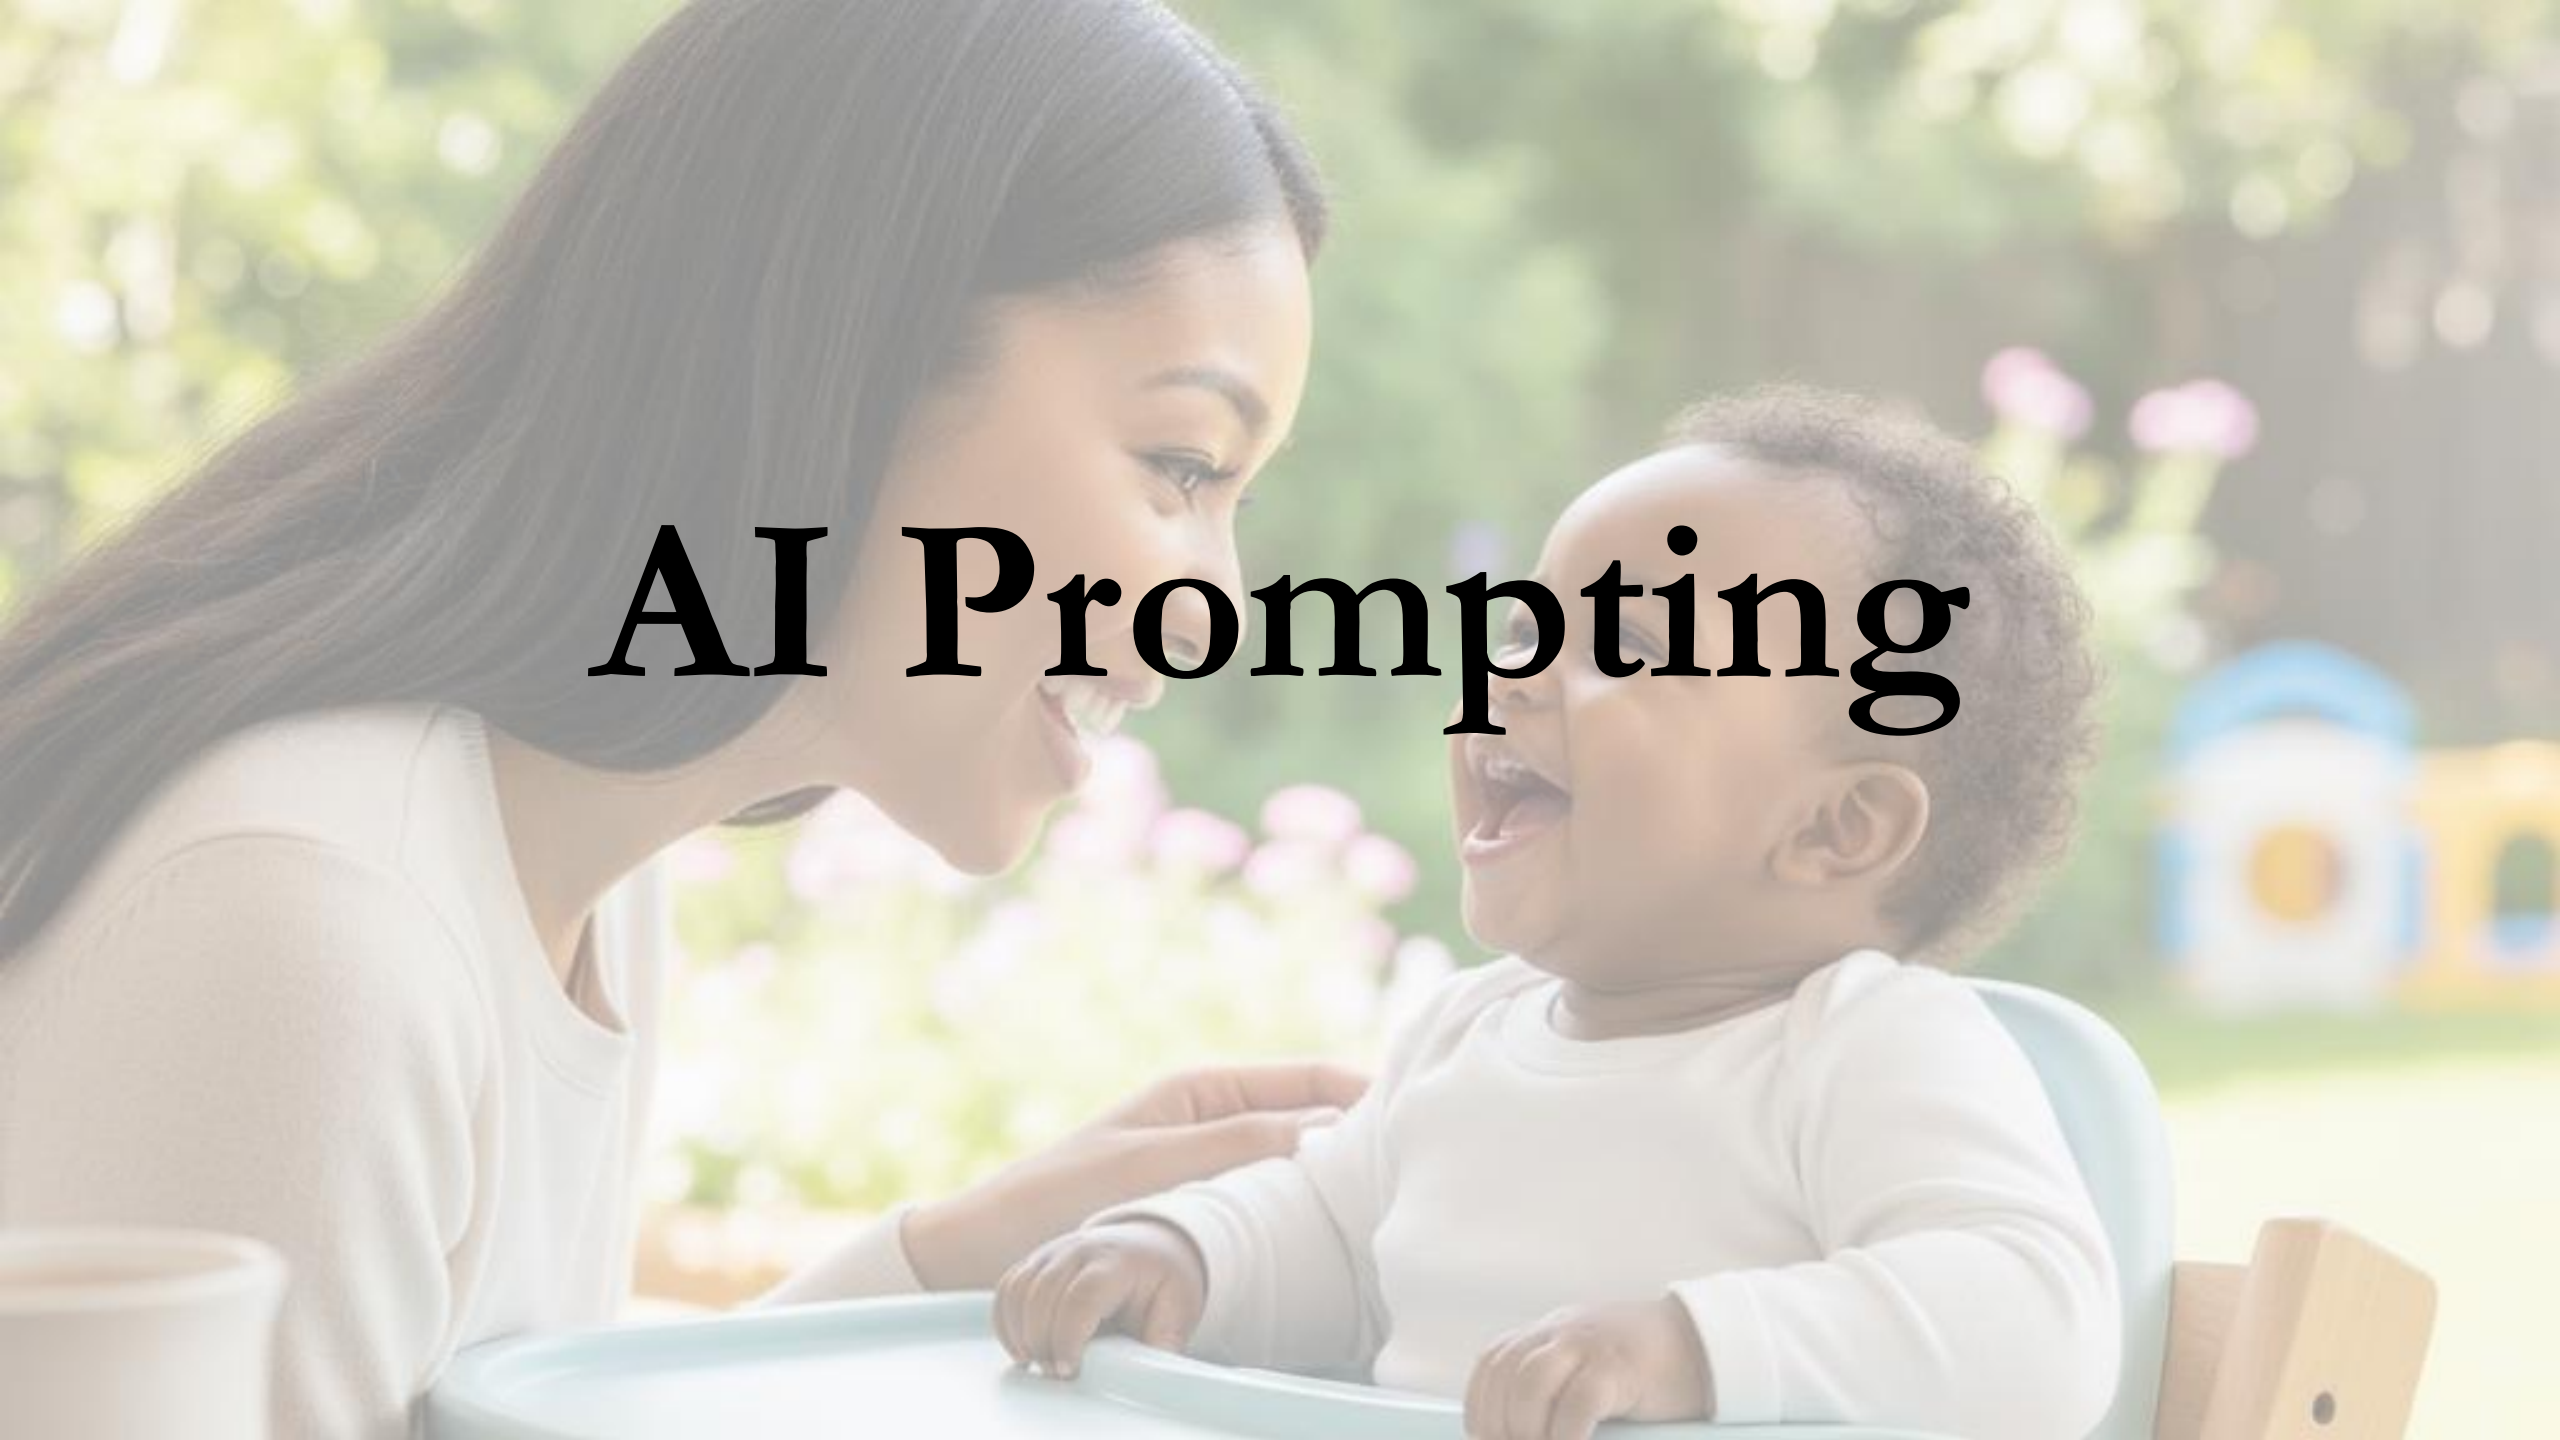

AI Prompting

The Challenge: "Garbage In, Garbage Out"

The "Blank Page" Problem

Without structure, AI models are forced to guess your intent. They act like a new intern who has memorized every document but has no real-world context or common sense.

The RTCF Solution

The **Role-Task-Context-Format** framework is a simple "cheat sheet" that provides the necessary constraints, turning generic outputs into actionable, professional results.

The RTCF Framework Components

Role

Who is the AI? Give it a job title (e.g., "Senior Policy Advisor") to prime the vocabulary.

Task

What do you want? Use strong action verbs (e.g., "Draft," "Summarize," "Analyze").

Context

Why/Who? Define the audience and constraints (e.g., "For busy executives").

Format

What does it look like? Specify structure (e.g., "Bulleted list," "Table").

Example: Policy Summary

❌ **Weak Prompt:**

"Summarize this report on water safety."

✅ **RTCF Prompt:**

- **Role:** Communications Director.
- **Task:** Summarize key findings.
- **Context:** Audience is general public; tone must be reassuring but urgent. No jargon.
- **Format:** Bulleted list + "What You Can Do" section.

Example: Operational Analysis

❌ **Weak Prompt:**

"Look at this data and tell me what it says."

✅ **RTCF Prompt:**

Role: Cynical Data Scientist looking for errors.

Task: Analyze monthly spending spreadsheet.

Context: We suspect travel expenses are too high this quarter. Flag anomalies.

Format: Markdown table highlighting top 5 suspicious entries.

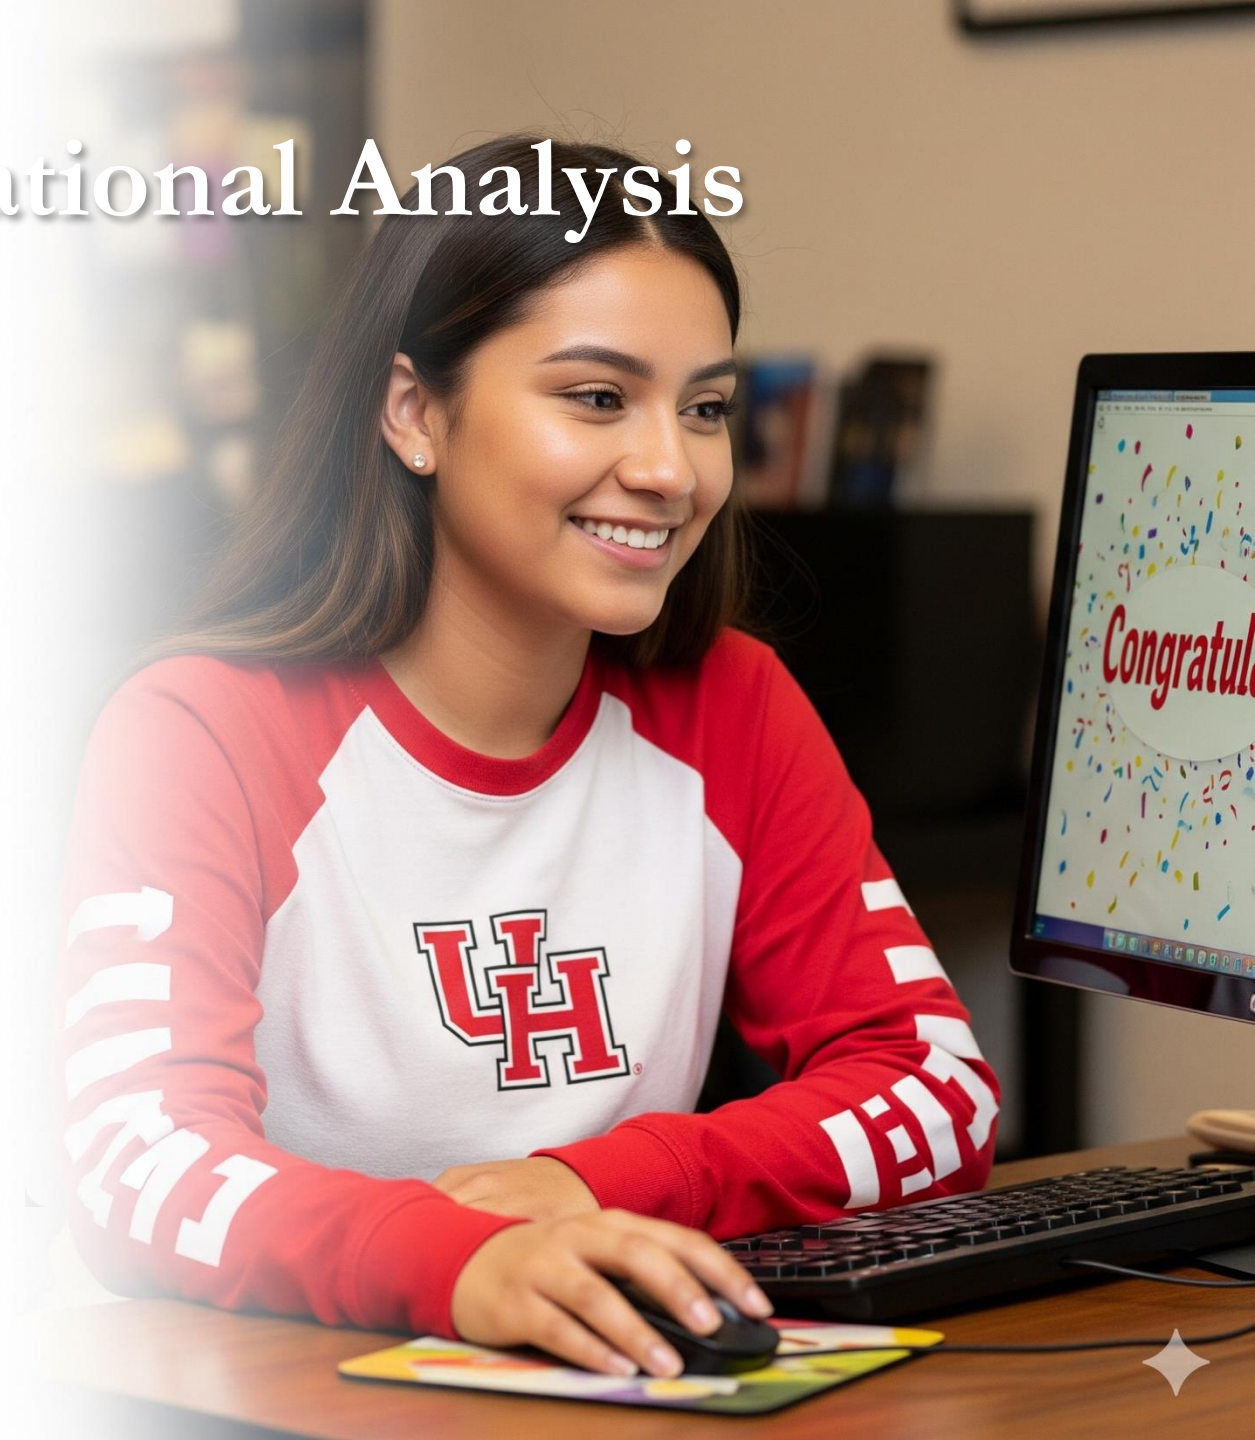

Example: The "Difficult Email"

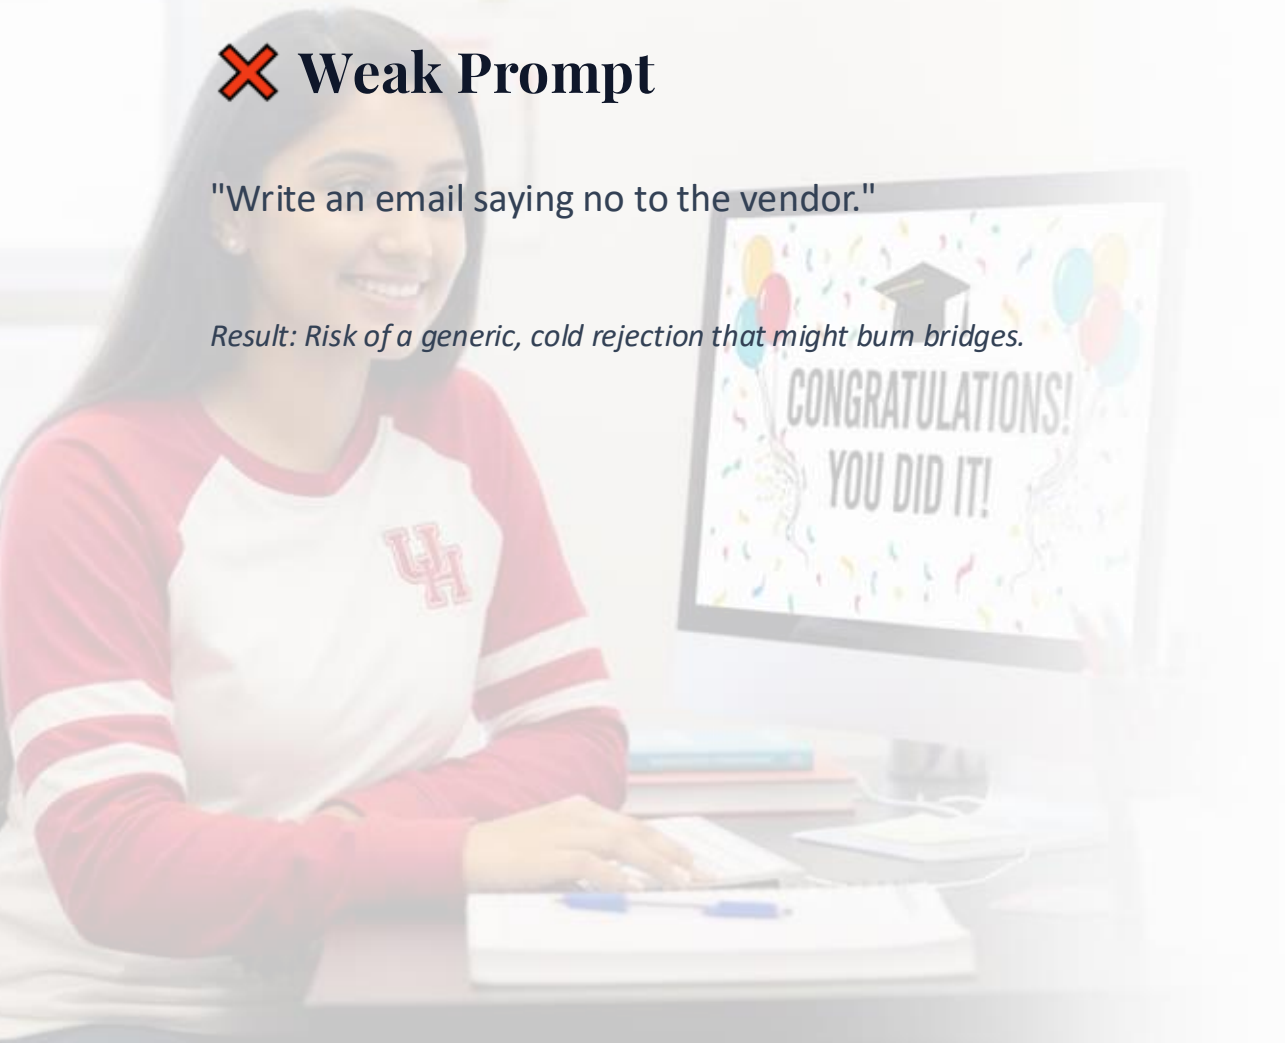

❌ Weak Prompt

"Write an email saying no to the vendor."

Result: Risk of a generic, cold rejection that might burn bridges.

✅ RTCF Prompt

- **Role:** Polite but firm Procurement Officer.
- **Task:** Draft rejection email.
- **Context:** Proposal was good but 20% over budget. We want to keep the door open for future work.
- **Format:** Professional email with warm opening and invitation to bid next year.

Summary Checklist (RTCF)

Component	The Question to Ask Yourself
Role	"Who should the AI pretend to be to do this best?"
Task	"What is the single specific action I need?"
Context	"Who is reading this? What does the AI *not* know?"
Format	"Do I want a table, a list, or a paragraph?"

IMPLEMENTING AI

RECOMENDATIONS

USE CASES

Domain	Use Case Example	The "Public Value"
Citizen Services	24/7 Intelligent Chatbots that handle Tier 1 queries (e.g., DMV hours, permit application status).	Reduces wait times; frees staff to handle complex casework.
Operations	Predictive Maintenance for fleets and public transit (e.g., predicting bus engine failure before it happens).	Saves tax dollars on repairs; improves service reliability.
Fraud & Compliance	Anomaly Detection in procurement or tax filings to flag irregular patterns.	Recovers revenue; protects public trust.
Policy & Analysis	Document Synthesis to summarize thousands of public comments on a proposed regulation.	Democratizes input; speeds up the policy cycle.

Risk & Governance: The ABC's of "How"

Government leaders face higher scrutiny than private sector CEOs. Your governance framework must address these unique risks.

A-B-C's

The How

Risk & Governance: The ABC's of "How"

A. Bias and Fairness

Risk: AI models trained on historical data often inherit historical inequities (e.g., a hiring algorithm rejecting candidates from certain zip codes).

Action: Mandate "Algorithmic Impact Assessments" before deployment to test for disparate impact on protected classes.

A-B-C's

The How

Risk & Governance: The ABC's of "How"

B. Data Privacy & Security

Risk: Government holds sensitive data (tax, health, justice). Generative AI tools can accidentally "leak" this data if staff paste sensitive info into public models.

Action: Enforce a "human-in-the-loop" policy and procure enterprise grade AI tools that guarantee data isolation (data is not used to train public models).

A-B-C's

The How

Risk & Governance: The ABC's of "How"

C. Hallucinations

Risk: Generative AI can confidently present false information as fact.

Action: Never use AI as the final approver for factual public communications. Establish a strict review process.

A-B-C's

The How

Risk & Governance: The ABC's of "How"

Implementation Roadmap for Leaders

Avoid "Pilot Purgatory" (where small tests never scale) by following this structured approach below:

Start Small, Think Big: Begin with low-risk internal tools (e.g., an internal HR bot) before launching public-facing AI.

Fix Your Data: AI is only as good as the data it is fed. If your agency's data is siloed or messy, AI will fail. Invest in data engineering first. "GIGO"

Upskill, Don't Just Replace: AI will shift roles, not necessarily eliminate them. Train policy analysts to be "AI Editors" staff who can prompt AI for drafts and then rigorously verify the output.

Procurement is Policy: Update RFPs to require vendors to disclose *training data sources* and *explainability features*. Don't buy what you can't audit.

Essential Questions to Ask Your Team

When presented with an AI proposal, ask these three questions:

- *"What data was this model trained on, and do we have the right to use it?"*
- *"If the AI gets it wrong, what is the worst-case scenario for a constituent?"*
- *"Is there a human in the loop to override the system if necessary?"*

RECAP - What is AI?

Artificial Intelligence (AI): A broad term for systems that perform tasks typically requiring human intelligence, such as recognizing patterns, interpreting speech, or making predictions.

AI | Machine Learning | Gen AI | “Black Box”

Artificial Intelligence (AI): A broad term for systems that perform tasks typically requiring human intelligence, such as recognizing patterns, interpreting speech, or making predictions.

The What?

Machine Learning (ML): The "engine" of modern AI. Instead of being explicitly programmed with rules (e.g., "If X, then Y"), ML systems learn rules by analyzing vast amounts of data.

Gov Relevance: ML is excellent for fraud detection, predicting infrastructure maintenance needs, or analyzing public health trends.

Generative AI (GenAI): A newer subset of AI (like ChatGPT or Claude) that can *create* new content—text, code, images, or summaries, based on patterns it learned during training.

Gov Relevance: Useful for drafting policy memos, summarizing constituent feedback from town halls, or modernizing legacy code.

The "Black Box" Problem: Advanced AI models often cannot explain *how* they reached a specific decision.

Leadership Note: In government, where due process is mandatory, you must be cautious about using "black box" AI for high-stakes decisions (e.g., denying benefits or sentencing).

Sending A Message

“Hey Siri, send “persons name” a voice message.”

Note: Default, will disappear after 2 minutes

“Hey Siri, send “persons name” a message.”

Emails

“Hey Siri, send an email to “persons name”

Note: *You will be prompted for for subject and message.*

“Hey Siri, show me all emails from *“persons name”*

“Hey Siri, show me all unread emails

Passwords

“Hey Siri, show me my password for “application name”

Note: *You will have to tap on the area to see the password*

The Weather

“Hey Siri, show me the weather for today”

Note: *The more specific you are, the more detailed data you will recv.*

“Hey Siri, show me the hourly weather/temperature/rain for “Carthage, Tx” today”

Translation

“Hey Siri, how do you say “where is the bathroom in French”

Note: *The more specific you are, the more detailed data you will recv.*

Making Phone Calls

“Hey Siri, call “person name” on speaker

Note: *You can also end the call if you activate Siri*

Finding People & Things

“Hey Siri, where is my “Mac Book Pro”

Note: *The more specific you are, the more detailed data you will recv. It will also locate airtags.*

“Hey Siri, where is my “person name”

“Hey Siri, where is my “car”

Note: *You must turn on this feature in maps (Show parked location) is toggled on*

Share Your Location

“Hey Siri, share my location with “Regina Jones”

“Hey Siri, where am I?”

Take A Screenshot

“Hey Siri, take a screenshot

Note: *Captures whatever is on your phone screen*

Basic Math

“Hey Siri, take a screen shot and send it to “persons name”

Note: *You can add a message also*

Time

“Hey Siri, how long until “time”

Note: *You can add a message also*

“Hey Siri, how many days until “date”

“Hey Siri, set a timer for “30 mins”

“Hey Siri, set a timer an alarm for “3:30 today”

Planning

“Hey Siri, am I free on Friday

Note: *You can add a message also*

“Hey Siri, am I free on Tuesday ay 10:30a

14

Setting Meetings and Reminders

“Hey Siri, create a meeting on this Friday and 10:30a

Note: *You can add a message also*

“Hey Siri, am I free on Tuesday ay 10:30a

Q&A

Prof. Kevin D. Jones
kevin@kevindjones.net