FacultyHack26

AI Paradigms & The Data Revolution

Understanding how machines learn and why the scale of modern data changed everything.

The Big Picture: AI vs. ML vs. DL

Defining the Hierarchy

Artificial Intelligence: The broad concept of machines acting "smart."

Machine Learning: A subset of AI where the computer improves at a task through experience without being explicitly programmed for every step.

Deep Learning: A specialized ML technique using "Neural Networks" inspired by the human brain to solve complex problems.

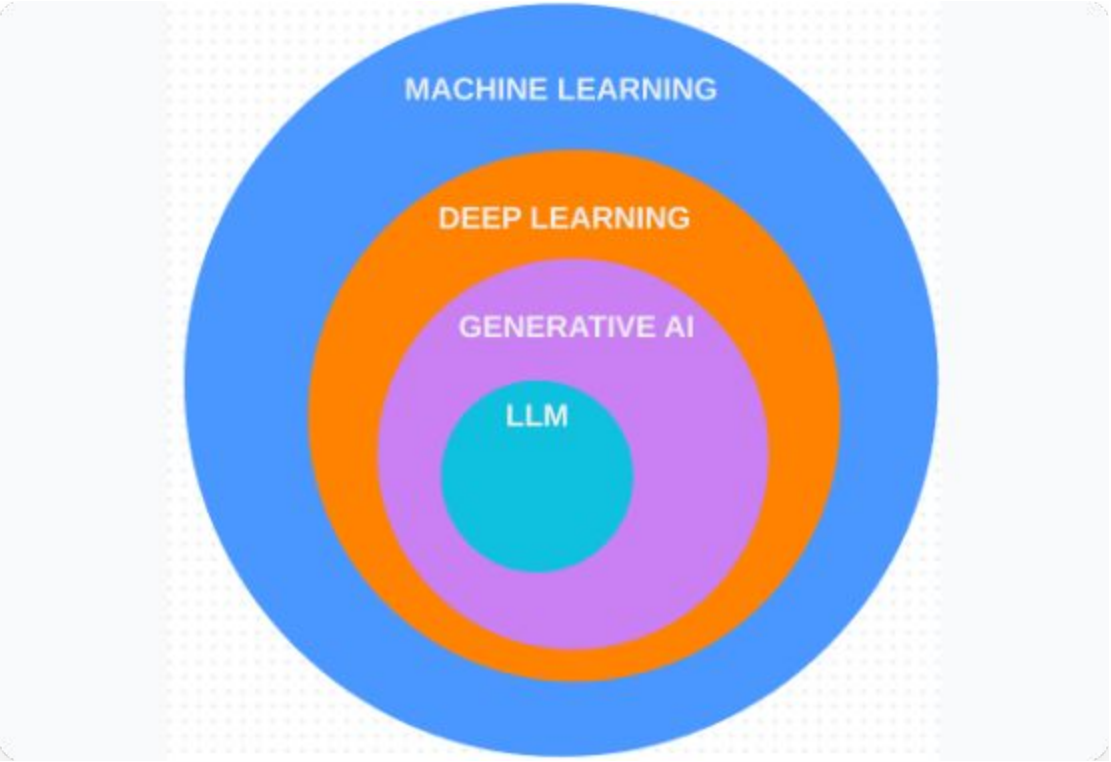

Supervised Learning: The Flashcard Way

The Analogy

Imagine a student studying with flashcards. On one side is the question (Input), and on the other is the correct answer (Label).
The "Teacher" provides thousands of labeled examples until the student recognizes the patterns.

Real-World Examples

- **Spam Detection:** Millions of emails marked as "Spam" or "Not Spam."
- **Medical Imaging:** X-rays labeled "Healthy" or "Tumor" by experts.

THINK ABOUT IT:
If the human "Teacher" makes a mistake and labels a cat as a dog on the flashcards, what will happen to the AI's learning?

Unsupervised: The Pattern Finder

The Analogy

Imagine a toddler with a big bin of mixed toys. No one tells them what the toys are, but they naturally start grouping them by color or shape.
The AI looks for hidden structures in data without being given any "correct" answers.

Real-World Examples

- **Customer Segments:** Grouping shoppers by their buying habits to find "The Bargain Hunters."
- **Anomaly Detection:** Finding a transaction that looks "weird" compared to your usual spending.

GROUP CHAT:
Can you think of a time a website recommended something to you that felt like it "knew" your taste, even if you never told it?

Reinforcement: The Puppy Trainer

Trial, Error, and Reward

The Analogy: Training a puppy. If it sits, it gets a treat (Reward). If it jumps, it gets nothing (Penalty).

The AI (Agent) interacts with an environment and learns a "Policy" to maximize its total rewards over time.

Example: Self-Driving Cars learning to navigate a street by receiving "points" for staying in lanes and "penalties" for hitting curbs.

Side-by-Side: ML Paradigms

Type	Input Data	Objective	Best Analogy
Supervised	Labeled (Questions + Answers)	Predict the next label	Teacher with Flashcards
Unsupervised	Unlabeled (Just Questions)	Find hidden patterns	Toddler sorting blocks
Reinforcement	Interaction (State + Action)	Maximize a Reward	Training a Puppy

The Data Revolution

Why did AI stay in research labs for 50 years and then suddenly take over the world? The answer is scale.

Exponential Growth of Data

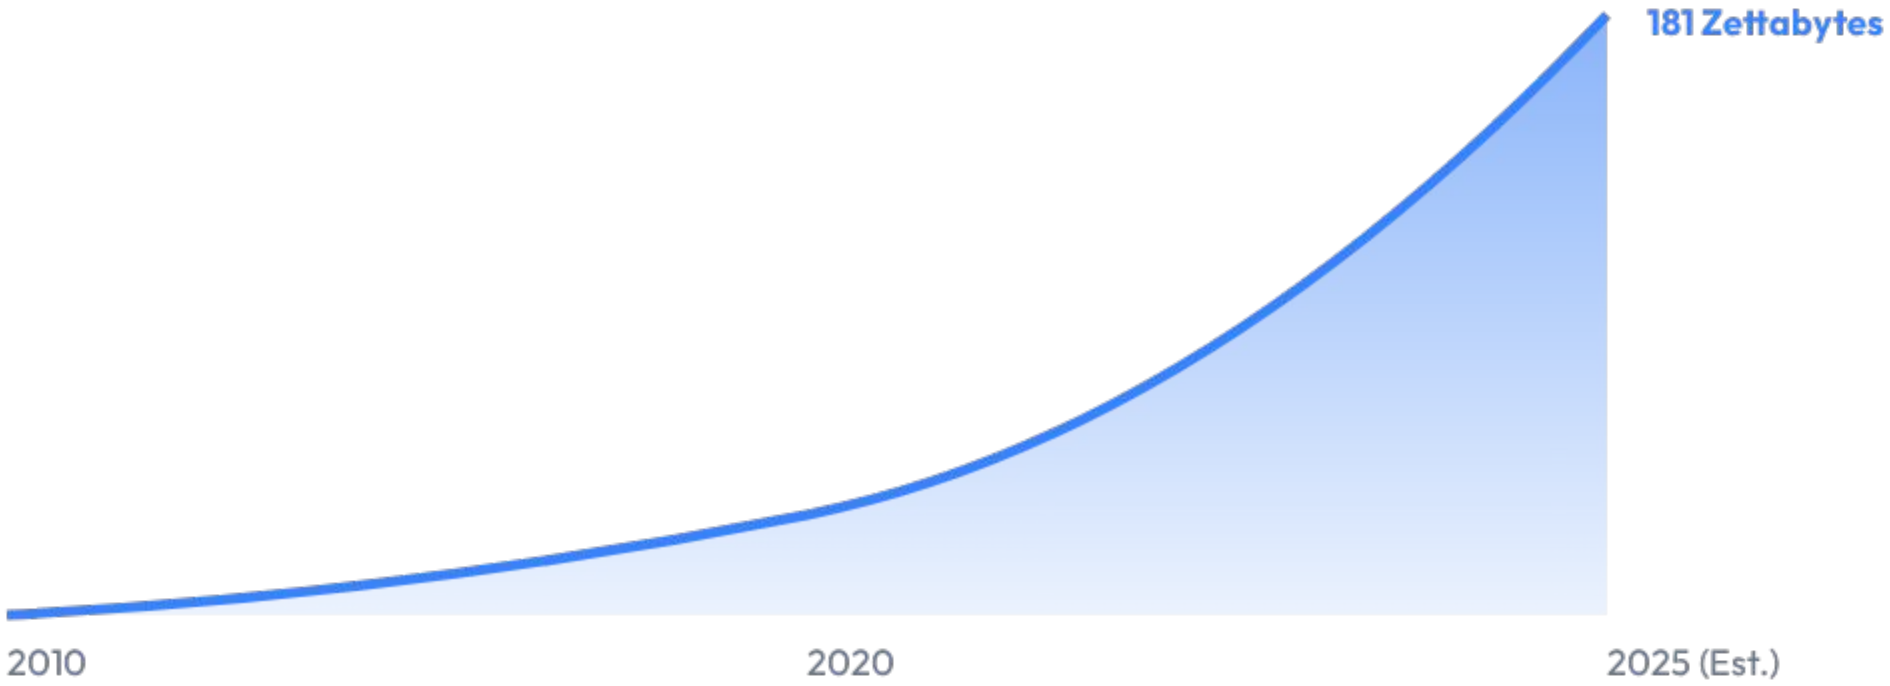

"Deep Learning models are like jet engines: they are incredibly powerful, but they need massive amounts of 'fuel' (data) to fly."

Why Ubiquity Drives Performance

Hardware Leap

GPUs (originally for gaming) allowed us to process massive data in parallel, making training fast.

Internet Everywhere

Billions of sensors, phones, and social posts provided the diverse examples AI needs.

Edge Cases

More data means the AI sees rare events (like a unicyclist on a highway) and learns to handle them.

A Stunning Reality

90%

Recent Data Surge

Roughly 90% of the world's total data has been created in just the last 2 years.

This "Ubiquity of Data" is what allowed models like ChatGPT to go from "interesting experiments" to "world-changing tools."

CRITICAL QUESTION:

If we run out of "new" human-made data to train on, will AI stop getting smarter? How might we solve that?

> **Algorithms are important, but data is the ground truth. Large-scale data didn't just improve AI; it changed what was possible.**
>
> — Random TACCster

Questions?