Fundamentals of AI/ML: From Turing Test to Transformers

Parikshit Pareek

Department of Electrical Engineering, Indian Institute of Technology Roorkee

July 18, 2025

What is Artificial Intelligence?

Definition: The ability of machines to perform tasks that typically require human intelligence

Key Components:

- Learning from experience
- Reasoning and problem-solving
- Understanding language
- Perception and recognition
- Planning and decision-making

Fun Fact: The term "Artificial Intelligence" was coined just to avoid using "cybernetics" because the researchers didn't want Norbert Wiener as their guru!

Alan Turing – The Father of AI (1912–1954)

Alan Turing (1912–1954)

Alan Turing – The Father of AI (1912–1954)

Alan Turing (1912–1954)

The Turing Machine (1936):

- Theoretical foundation for all modern computers
- Universal computation concept
- Algorithm formalization

Alan Turing – The Father of AI (1912–1954)

Alan Turing (1912–1954)

The Turing Machine (1936):

- Theoretical foundation for all modern computers
- Universal computation concept
- Algorithm formalization

The Turing Test (1950):

- "Can machines think?"
- Imitation Game concept
- Still relevant today for evaluating AI

Alan Turing – The Father of AI (1912–1954)

Alan Turing (1912–1954)

The Turing Machine (1936):

- Theoretical foundation for all modern computers
- Universal computation concept
- Algorithm formalization

The Turing Test (1950):

- "Can machines think?"
- Imitation Game concept
- Still relevant today for evaluating AI

Wartime Contributions:

- Enigma code breaking at Bletchley Park
- Shortened WWII by an estimated 2–4 years

The Birth of AI - Dartmouth Conference (1956)

The Constitutional Convention of AI– Organizers:

- John McCarthy (coined "Artificial Intelligence"; pioneer of Lisp and AI research)
- Marvin Minsky (cognitive scientist; co-founder of MIT AI Lab)
- Claude Shannon (father of information theory)
- Nathaniel Rochester (IBM engineer; designer of the IBM 701)

Goals:

- "Every aspect of learning can be precisely described so that a machine can simulate it"
- 6-8 week brainstorming session
- Foundation of AI as a research field

Outcome:

Birth of AI optimism - "Machines will be capable of doing any work a man can do" within 20 years

Early AI Achievements (1950s-1960s)

ELIZA (1966):

- First chatbot by Joseph Weizenbaum
- Simulated psychotherapist
- Demonstrated how simple pattern matching could fool users

Logic Theorist (1956):

- First AI program
- Proved mathematical theorems
- Written during Dartmouth Conference

General Problem Solver (1957):

- Herbert Simon and Allen Newell
- First general-purpose problem-solving program

The First AI Winter (1974-1980)

What Went Wrong?

- Overpromising and underdelivering
- Limited computational power
- Lack of real-world applications
- The Lighthill Report (1973) criticized AI research

Key Problems:

- Combinatorial explosion
- Frame problem
- Common sense reasoning
- Limited memory and processing power

Funding Impact:

- DARPA cut AI funding
- British government reduced support
- Academic positions disappeared

Lesson: Hype cycles are dangerous in AI!

The Expert Systems Boom (1980-1987)

What are Expert Systems?

- Programs that capture expert knowledge in specific domains
- Rule-based reasoning
- If-then logic structures

Success Stories:

- **MYCIN (1972):** Medical diagnosis
- **XCON/R1 (1980):** Saved DEC $40 million annually
- **DENDRAL (1965):** Chemical analysis

Why They Worked:

- Narrow, well-defined domains
- Clear rules and logic
- Practical business value

The Second AI Winter (1987-1993)

The Expert Systems Bubble Burst:

- Expensive to maintain
- Brittle and inflexible
- Couldn't handle uncertainty
- Knowledge acquisition bottleneck

Economic Factors:

- LISP machine market collapsed
- Japan's Fifth Generation project failed
- Desktop PCs became more powerful than LISP machines

Silver Lining: This winter led to more realistic expectations and better methodologies

Computing Revolution Enables AI Renaissance

Moore's Law Impact:

- Computing power doubles every 2 years
- Enabled complex algorithms
- Made large-scale data processing possible

GPU Revolution:

- Originally for graphics
- Perfect for parallel computation
- Neural networks need massive parallelization
- NVIDIA becomes AI kingmaker

The Internet:

- Massive datasets available
- Distributed computing
- Real-world testing platforms

Timeline of AI

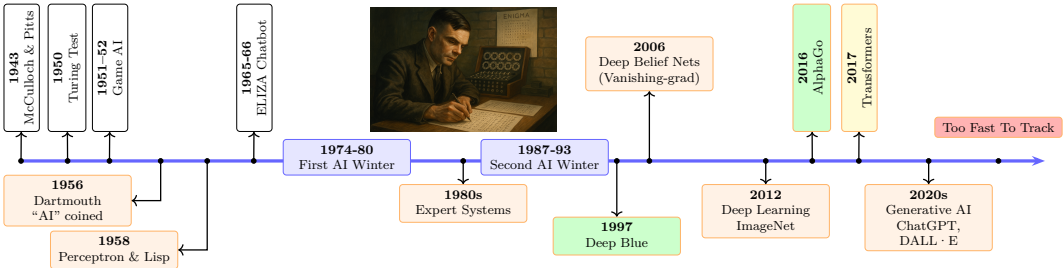

GPU vs CPU for AI

CPU:

- Few powerful cores
- Sequential processing
- Complex instruction sets
- Good for general computing

GPU:

- Thousands of simple cores
- Parallel processing
- Perfect for matrix operations
- 10x faster for neural networks

Key Insight: AI workloads are embarrassingly parallel - perfect for GPUs!

AI, Machine Learning, and Deep Learning

ANY COMPUTER PROGRAM WHICH SHOWS INTELLIGENCE LIKE CHESS MOVES

Fun Analogy:
AI is like "intelligence"
ML is like "learning from examples,"
DL is like "learning REALLY hard & LOTS OF examples!"

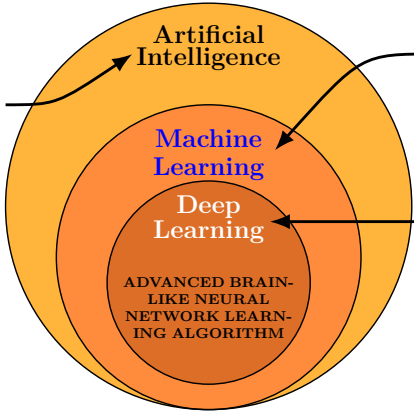

TRADITIONAL STATISTICAL PREDICTION MODELS WHICH ALLOWS FOR OPTIMIZATION LIKE FORECASTING

LARGER MODELS WHICH FEEDS ON LOADS OF DATA & TRAINS ON LOTS OF COMPUTE LIKE CHATGPT

Machine Learning Types

	Supervised Learning	**Unsupervised Learning**	**Reinforcement Learning**
Description	Trained on Labeled Data	Trained on Unlabeled Data	Learn by Interacting with Environment
Goal	Make Predictions	Find Patterns and Structure	Maximize Reward
Examples	Spam Detection Image Classification	Clustering Dimensionality Reduction	Game Playing Robotics LLMs

The Deep Learning Revolution (2006-2012)

Geoffrey Hinton's Comeback:

- Coined "Deep Learning" in 2006
- Solved the vanishing gradient problem
- Layer-by-layer pre-training

AlexNet (2012):

- Won ImageNet competition
- 84% $\rightarrow$ 15% error rate improvement
- Proved deep learning's potential

Key Enablers:

- Big data (ImageNet, etc.)
- GPU computing
- Better algorithms (ReLU, dropout, batch normalization)

AI Compute Growth

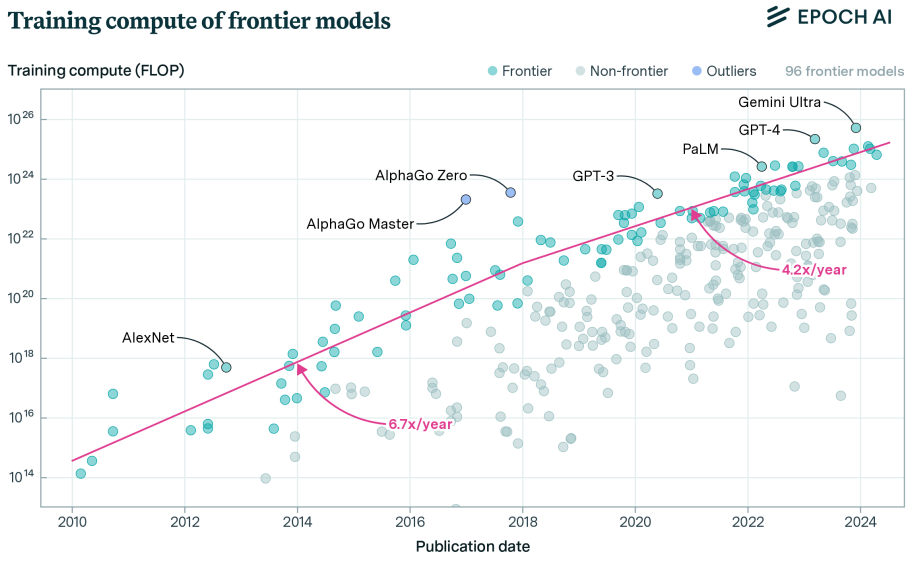

AI Today - Capabilities and Limitations

What AI Can Do Well:

- Pattern recognition
- Language translation
- Code generation
- Creative writing
- Game playing
- Image generation

What AI Still Struggles With:

- Common sense reasoning
- Causal understanding
- Long-term planning
- Robustness and safety
- Understanding vs. pattern matching

Future Directions and Challenges

Technical Challenges:

- Artificial General Intelligence (AGI)
- Explainable AI
- Few-shot learning
- Multimodal understanding
- Energy efficiency

Societal Challenges:

- Job displacement
- Bias and fairness
- Privacy and surveillance
- Autonomous weapons
- Misinformation

Questions!