EEC 351 Fundamentals of AI/ML: From the Turing Test to World Models

Lecture 1 — A Short History, and What It Teaches Us

Parikshit Pareek

Department of Electrical Engineering
Indian Institute of Technology Roorkee

Autumn 2026–27

Current Stage, in five acts

1. **Birth** — Turing's question, and the Dartmouth summer that named a field

Current Stage, in five acts

1. **Birth** — Turing's question, and the Dartmouth summer that named a field
2. **Winters** — two spectacular collapses, and what actually froze

Current Stage, in five acts

1. **Birth** — Turing's question, and the Dartmouth summer that named a field
2. **Winters** — two spectacular collapses, and what actually froze
3. **Revival** — the hardware and data that thawed everything

Current Stage, in five acts

1. **Birth** — Turing's question, and the Dartmouth summer that named a field
2. **Winters** — two spectacular collapses, and what actually froze
3. **Revival** — the hardware and data that thawed everything
4. **Now** — deep learning, transformers, and the road to world models

Current Stage, in five acts

1. **Birth** — Turing's question, and the Dartmouth summer that named a field
2. **Winters** — two spectacular collapses, and what actually froze
3. **Revival** — the hardware and data that thawed everything
4. **Now** — deep learning, transformers, and the road to world models
5. **The moral** — what 70 years of boom and bust says about *this course*

Current Stage, in five acts

1. **Birth** — Turing's question, and the Dartmouth summer that named a field
2. **Winters** — two spectacular collapses, and what actually froze
3. **Revival** — the hardware and data that thawed everything
4. **Now** — deep learning, transformers, and the road to world models
5. **The moral** — what 70 years of boom and bust says about *this course*

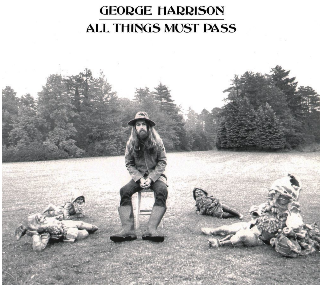

What is Artificial Intelligence?

Working definition: the ability of machines to perform tasks that typically require human intelligence.

What is Artificial Intelligence?

Working definition: the ability of machines to perform tasks that typically require human intelligence.

The classic checklist:

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

What is Artificial Intelligence?

Working definition: the ability of machines to perform tasks that typically require human intelligence.

The classic checklist:

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

Naming trivia

McCarthy chose the name "Artificial Intelligence" partly to keep the new field at arm's length from *cybernetics* — and from having Norbert Wiener as its guru. Fields, like startups, are sometimes named for who they want to avoid.

Alan Turing — the question before the field (1912–1954)

Alan Turing (1912–1954)

The Turing Machine (1936):

- Theoretical foundation of all modern computers
- Universal computation; the formalisation of "algorithm"

Alan Turing — the question before the field (1912–1954)

Alan Turing (1912–1954)

The Turing Machine (1936):

- Theoretical foundation of all modern computers
- Universal computation; the formalisation of "algorithm"

The Turing Test (1950):

- "Can machines think?" recast as the *Imitation Game*
- Still shaping how we argue about LLMs today

Alan Turing — the question before the field (1912–1954)

Alan Turing (1912–1954)

The Turing Machine (1936):

- Theoretical foundation of all modern computers
- Universal computation; the formalisation of "algorithm"

The Turing Test (1950):

- "Can machines think?" recast as the *Imitation Game*
- Still shaping how we argue about LLMs today

Wartime work:

- Enigma code-breaking at Bletchley Park
- Estimated (by some historians) to have shortened WWII by years

The birth certificate of AI — Dartmouth, summer 1956

The proposal's authors:

- John McCarthy — coined "Artificial Intelligence"; later created Lisp
- Marvin Minsky — co-founder of the MIT AI Lab
- Claude Shannon — father of information theory
- Nathaniel Rochester — IBM engineer; designer of the IBM 701

The birth certificate of AI — Dartmouth, summer 1956

The proposal's authors:

- John McCarthy — coined "Artificial Intelligence"; later created Lisp
- Marvin Minsky — co-founder of the MIT AI Lab
- Claude Shannon — father of information theory
- Nathaniel Rochester — IBM engineer; designer of the IBM 701

The famous sentence from the proposal:

> *"Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."*

The birth certificate of AI — Dartmouth, summer 1956

The proposal's authors:

- John McCarthy — coined "Artificial Intelligence"; later created Lisp
- Marvin Minsky — co-founder of the MIT AI Lab
- Claude Shannon — father of information theory
- Nathaniel Rochester — IBM engineer; designer of the IBM 701

The famous sentence from the proposal:

> *"Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."*

Outcome

A 6–8 week brainstorming summer founded a research field — and its first hype cycle: confident predictions of human-level machines "within 20 years."

Early wins (1955–1966): three programs, three lessons

Logic Theorist (1956)

- Newell & Simon; presented at Dartmouth
- Proved theorems from *Principia Mathematica*
- Lesson: reasoning can be *search*

Early wins (1955–1966): three programs, three lessons

Logic Theorist (1956)

- Newell & Simon; presented at Dartmouth
- Proved theorems from *Principia Mathematica*
- Lesson: reasoning can be *search*

General Problem Solver (1957)

- Newell & Simon again
- First general-purpose problem solver
- Lesson: generality is *hard* — it worked only on toy problems

Early wins (1955–1966): three programs, three lessons

Logic Theorist (1956)

- Newell & Simon; presented at Dartmouth
- Proved theorems from *Principia Mathematica*
- Lesson: reasoning can be *search*

General Problem Solver (1957)

- Newell & Simon again
- First general-purpose problem solver
- Lesson: generality is *hard* — it worked only on toy problems

ELIZA (1966)

- Weizenbaum's chatbot "psychotherapist"
- Simple pattern matching — yet users confided in it
- Lesson: humans *project* understanding. Keep this in mind for every LLM debate you will ever have.

The First AI Winter (1974–1980)

What went wrong:

- Overpromising, underdelivering
- Combinatorial explosion — toy demos didn't scale
- Common-sense reasoning: nowhere in sight
- The Lighthill Report (1973) said so, in public

The First AI Winter (1974–1980)

What went wrong:

- Overpromising, underdelivering
- Combinatorial explosion — toy demos didn't scale
- Common-sense reasoning: nowhere in sight
- The Lighthill Report (1973) said so, in public

What it cost:

- DARPA cut AI funding
- British support collapsed post-Lighthill
- Academic positions disappeared

The lesson

Hype is a loan. Winters are how the field pays it back, with interest.

The Expert Systems Boom (1980–1987)

The idea: capture an expert's knowledge as *if–then rules*; reason over the rule base.

The Expert Systems Boom (1980–1987)

The idea: capture an expert's knowledge as *if–then rules*; reason over the rule base.

Genuine successes:

- **DENDRAL (1965–):** chemical structure analysis — the prototype
- **MYCIN (1970s):** medical diagnosis; matched specialists in trials
- **XCON/R1 (1980):** configured DEC computers; reportedly saved $40M a year

The Expert Systems Boom (1980–1987)

The idea: capture an expert's knowledge as *if–then rules*; reason over the rule base.

Genuine successes:

- **DENDRAL (1965–):** chemical structure analysis — the prototype
- **MYCIN (1970s):** medical diagnosis; matched specialists in trials
- **XCON/R1 (1980):** configured DEC computers; reportedly saved $40M a year

Why they worked: narrow, well-defined domains; explicit rules; clear business value. **Remember this pattern** — *narrow and well-specified* is where AI keeps succeeding first, in every era.

The Second AI Winter (1987–1993)

The bubble burst because expert systems were:

- Expensive to build *and* maintain
- Brittle — no graceful handling of uncertainty
- Bottlenecked on knowledge acquisition: rules had to be typed in by hand

The Second AI Winter (1987–1993)

The bubble burst because expert systems were:

- Expensive to build *and* maintain
- Brittle — no graceful handling of uncertainty
- Bottlenecked on knowledge acquisition: rules had to be typed in by hand

The economics finished the job:

- LISP machine market collapsed
- Japan's Fifth Generation project fell short
- Desktop PCs overtook specialised hardware

The Second AI Winter (1987–1993)

The bubble burst because expert systems were:

- Expensive to build *and* maintain
- Brittle — no graceful handling of uncertainty
- Bottlenecked on knowledge acquisition: rules had to be typed in by hand

The economics finished the job:

- LISP machine market collapsed
- Japan's Fifth Generation project fell short
- Desktop PCs overtook specialised hardware

Silver lining

The field emerged humbler, more mathematical, and more empirical — the culture that would make machine *learning* respectable. (Notice what failed: typing knowledge in. Notice what would win: *learning it from data*.)

Seventy years in one curve

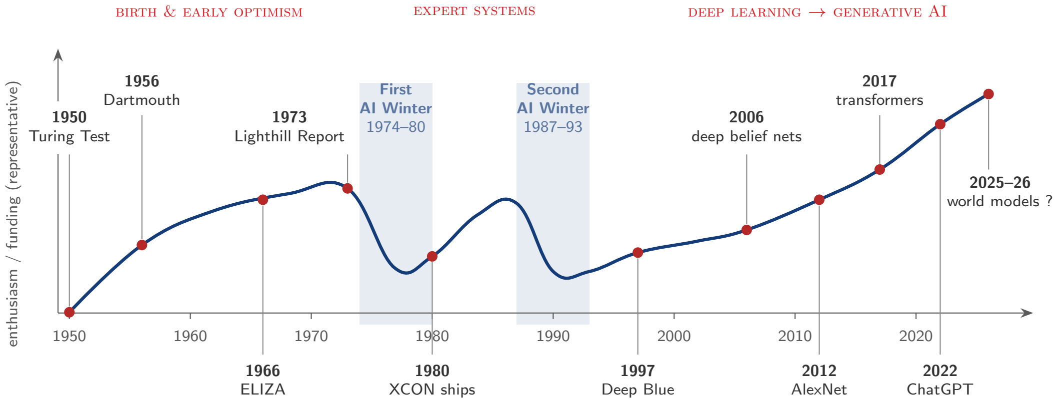

The curve is schematic with correct dates

What thawed the winters: three exponential gifts

Moore's Law

- Transistor counts doubling every $\sim$2 years
- Yesterday's intractable becomes today's homework

What thawed the winters: three exponential gifts

Moore's Law

- Transistor counts doubling every $\sim$2 years
- Yesterday's intractable becomes today's homework

The GPU accident

- Built for graphics; perfect for parallel matrix maths
- Neural networks are *made* of matrix maths
- NVIDIA becomes the field's kingmaker

What thawed the winters: three exponential gifts

Moore's Law

- Transistor counts doubling every $\sim$2 years
- Yesterday's intractable becomes today's homework

The GPU accident

- Built for graphics; perfect for parallel matrix maths
- Neural networks are *made* of matrix maths
- NVIDIA becomes the field's kingmaker

The Internet

- Datasets of unprecedented scale
- Distributed computing
- Billions of users as a testing ground

What thawed the winters: three exponential gifts

Moore's Law

- Transistor counts doubling every $\sim$2 years
- Yesterday's intractable becomes today's homework

The GPU accident

- Built for graphics; perfect for parallel matrix maths
- Neural networks are *made* of matrix maths
- NVIDIA becomes the field's kingmaker

The Internet

- Datasets of unprecedented scale
- Distributed computing
- Billions of users as a testing ground

None of these were built for AI. AI was the beneficiary.

GPU vs CPU: why the accident mattered

CPU:

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

GPU:

- Thousands of simple cores
- Massively parallel processing
- Built for matrix operations
- Order(s) of magnitude faster for neural network training

GPU vs CPU: why the accident mattered

CPU:

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

GPU:

- Thousands of simple cores
- Massively parallel processing
- Built for matrix operations
- Order(s) of magnitude faster for neural network training

Key insight — and a course preview

AI workloads are *embarrassingly parallel*: mostly enormous matrix–vector products. When we say in the Linear Algebra review that "a model is matrices," this slide is the hardware receipt.

AI ⊃ ML ⊃ DL ⊃ Generative AI

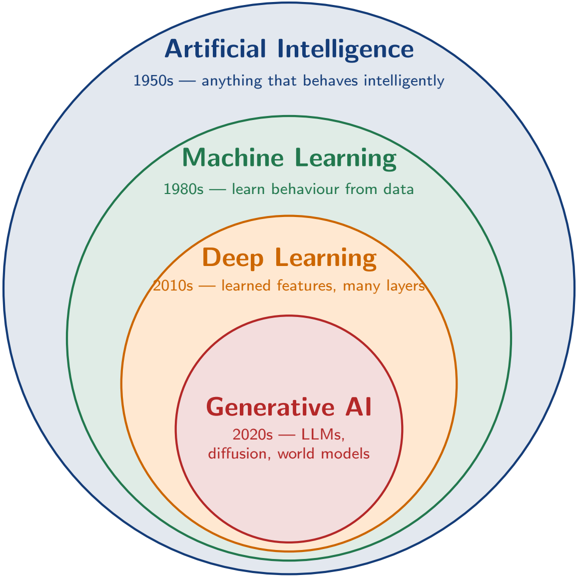

- **AI**: the ambition — rule-based systems, search, planning all count

AI ⊃ ML ⊃ DL ⊃ Generative AI

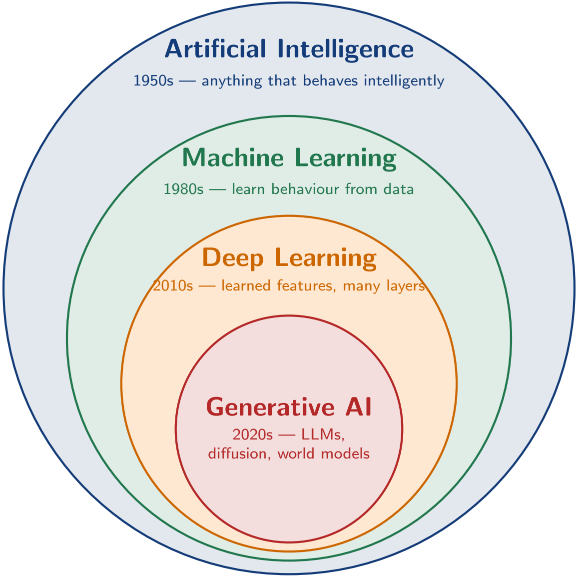

- **AI**: the ambition — rule-based systems, search, planning all count
- **ML**: the strategy — stop programming behaviour, *learn* it from data

AI ⊃ ML ⊃ DL ⊃ Generative AI

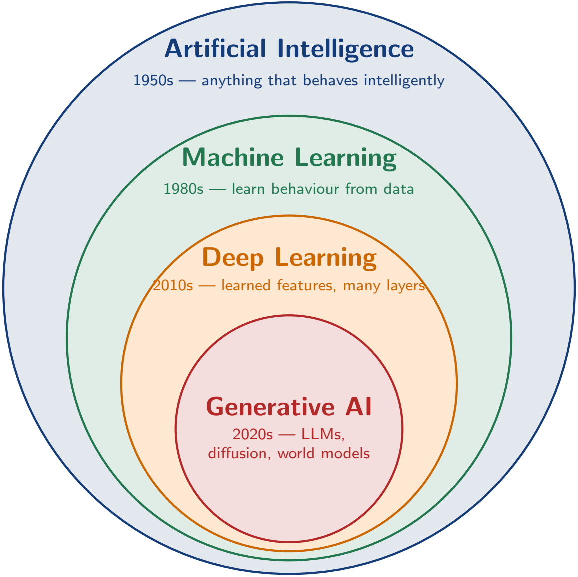

- **AI**: the ambition — rule-based systems, search, planning all count
- **ML**: the strategy — stop programming behaviour, *learn* it from data
- **DL**: the architecture — stop hand-crafting features, learn those too

AI ⊃ ML ⊃ DL ⊃ Generative AI

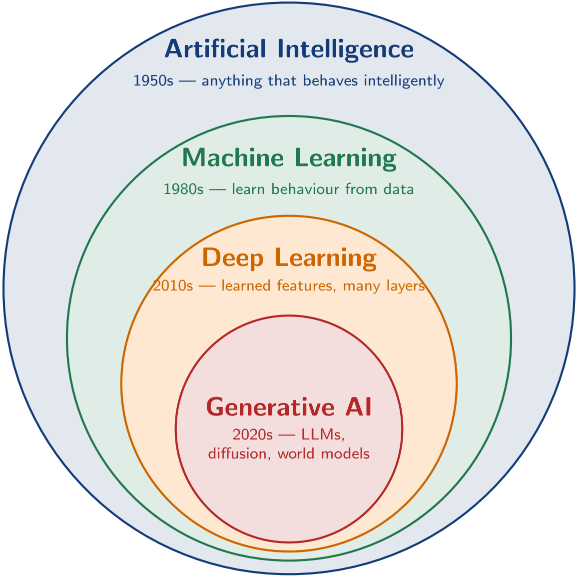

- **AI**: the ambition — rule-based systems, search, planning all count
- **ML**: the strategy — stop programming behaviour, *learn* it from data
- **DL**: the architecture — stop hand-crafting features, learn those too
- **Generative AI**: the current frontier — learn to *produce* language, images, video, actions

AI ⊃ ML ⊃ DL ⊃ Generative AI

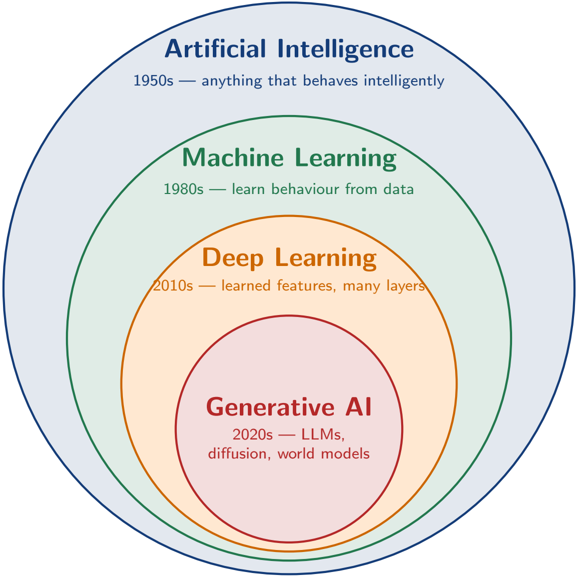

- **AI**: the ambition — rule-based systems, search, planning all count
- **ML**: the strategy — stop programming behaviour, *learn* it from data
- **DL**: the architecture — stop hand-crafting features, learn those too
- **Generative AI**: the current frontier — learn to *produce* language, images, video, actions

Each ring keeps the outer ring's goal and automates one more layer of human effort. This course lives mostly in rings 2–3 — deliberately: that is where the transferable mathematics is.

The Deep Learning Revolution (2006–2012)

The revival (2006):

- Hinton's deep belief networks showed deep architectures *could* be trained (layer-by-layer pre-training), reigniting the field under the banner "deep learning"

The Deep Learning Revolution (2006–2012)

The revival (2006):

- Hinton's deep belief networks showed deep architectures *could* be trained (layer-by-layer pre-training), reigniting the field under the banner "deep learning"

The moment everyone remembers — AlexNet (2012):

- Won the ImageNet challenge: top-5 error **26.2%** $\to$ **15.3%** — a jump the field had never seen
- Trained on two consumer GPUs; the "accident" pays off

The Deep Learning Revolution (2006–2012)

The revival (2006):

- Hinton's deep belief networks showed deep architectures *could* be trained (layer-by-layer pre-training), reigniting the field under the banner "deep learning"

The moment everyone remembers — AlexNet (2012):

- Won the ImageNet challenge: top-5 error **26.2%** $\rightarrow$ **15.3%** — a jump the field had never seen
- Trained on two consumer GPUs; the "accident" pays off

The three enablers (memorise this triple):

- **Data** (ImageNet's 1.2M labelled images)
- **Compute** (GPUs)
- **Algorithms** (ReLU, dropout, better initialisation)

2017 → 2026: from Attention to World Models

- **2017 —** "Attention Is All You Need": the **transformer**; attention scores token pairs by inner products $QK^\top$ (Linear Algebra, again)

2017 $\rightarrow$ 2026: from Attention to World Models

- **2017** — "Attention Is All You Need": the **transformer**; attention scores token pairs by inner products $QK^\top$ (Linear Algebra, again)
- **2018–20** — BERT; GPT-2; **GPT-3** (175B parameters): *scaling laws* — more data, parameters, compute $\Rightarrow$ predictably better models

2017 $\to$ 2026: from Attention to World Models

- **2017** — “Attention Is All You Need”: the **transformer**; attention scores token pairs by inner products $QK^\top$ (Linear Algebra, again)
- **2018–20** — BERT; GPT-2; **GPT-3** (175B parameters): *scaling laws* — more data, parameters, compute $\Rightarrow$ predictably better models
- **Nov 2022** — **ChatGPT**: 100M users in two months; AI becomes a consumer product

2017 $\rightarrow$ 2026: from Attention to World Models

- **2017** — "Attention Is All You Need": the **transformer**; attention scores token pairs by inner products $QK^\top$ (Linear Algebra, again)
- **2018–20** — BERT; GPT-2; **GPT-3** (175B parameters): *scaling laws* — more data, parameters, compute $\Rightarrow$ predictably better models
- **Nov 2022** — **ChatGPT**: 100M users in two months; AI becomes a consumer product
- **2023–25** — multimodal models; then *reasoning* models that spend compute at answer time thinking step-by-step

2017 $\rightarrow$ 2026: from Attention to World Models

- **2017** — "Attention Is All You Need": the **transformer**; attention scores token pairs by inner products $QK^\top$ (Linear Algebra, again)
- **2018–20** — BERT; GPT-2; **GPT-3** (175B parameters): *scaling laws* — more data, parameters, compute $\Rightarrow$ predictably better models
- **Nov 2022** — **ChatGPT**: 100M users in two months; AI becomes a consumer product
- **2023–25** — multimodal models; then *reasoning* models that spend compute at answer time thinking step-by-step
- **2025–26** — the **world-model** turn: systems that learn predictive models of environments (video, robotics) rather than only text — a topic with a starring role later in *this course* (energy-based modelling unit)

2017 $\rightarrow$ 2026: from Attention to World Models

- **2017** — "Attention Is All You Need": the **transformer**; attention scores token pairs by inner products $QK^\top$ (Linear Algebra, again)
- **2018–20** — BERT; GPT-2; **GPT-3** (175B parameters): *scaling laws* — more data, parameters, compute $\Rightarrow$ predictably better models
- **Nov 2022** — **ChatGPT**: 100M users in two months; AI becomes a consumer product
- **2023–25** — multimodal models; then *reasoning* models that spend compute at answer time thinking step-by-step
- **2025–26** — the **world-model** turn: systems that learn predictive models of environments (video, robotics) rather than only text — a topic with a starring role later in *this course* (energy-based modelling unit)

The pattern to notice

Every jump on this slide is the same three ingredients as 2012 — data, compute, algorithms — at ever larger scale.

The compute curve behind it all

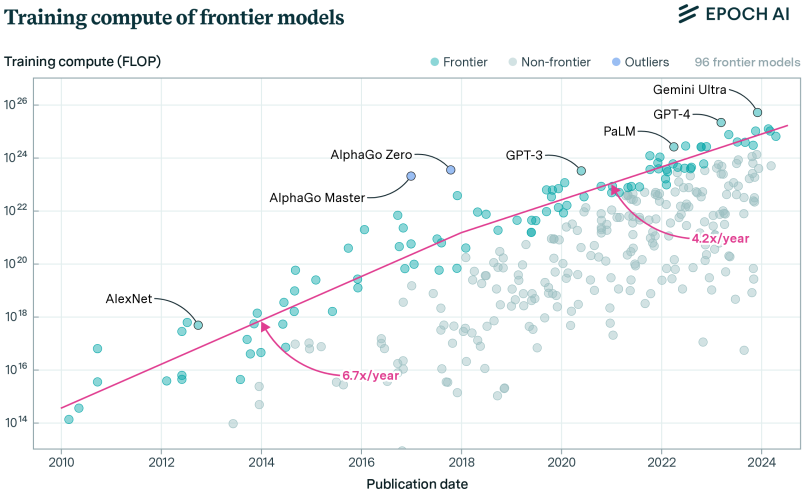

The compute curve behind it all

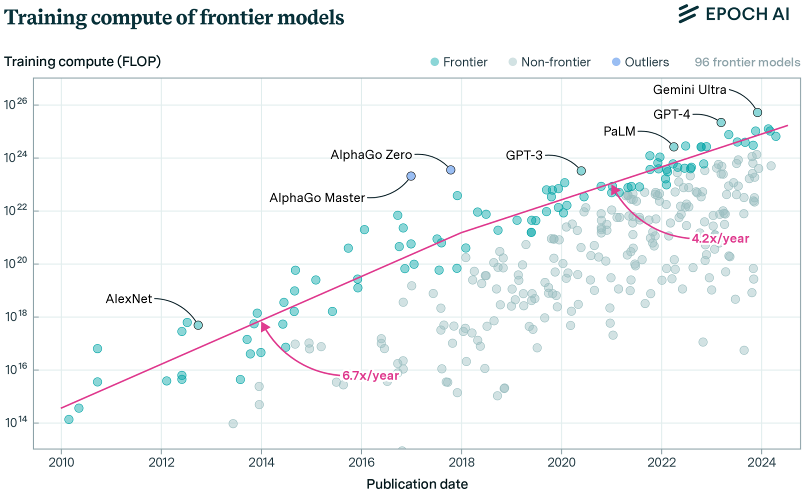

Notice the year graph is ending!

The three (and a half) types of machine learning

	Supervised	**Unsupervised / Self-supervised**	**Reinforcement**
Learns from	Labelled data	Unlabelled data (labels manufactured from the data itself)	Interaction with an environment
Goal	Predict labels	Find structure; learn representations	Maximise long-run reward
Examples	Spam detection Image classification	Clustering, PCA **LLM pretraining**	Game playing Robotics **RLHF fine-tuning**

The three (and a half) types of machine learning

	Supervised	**Unsupervised / Self-supervised**	**Reinforcement**
Learns from	Labelled data	Unlabelled data (labels manufactured from the data itself)	Interaction with an environment
Goal	Predict labels	Find structure; learn representations	Maximise long-run reward
Examples	Spam detection Image classification	Clustering, PCA **LLM pretraining**	Game playing Robotics **RLHF fine-tuning**

Where do LLMs sit?

Mixed Bag: *self-supervised* pretraining (predict the next token — the labels are free) + *reinforcement* fine-tuning (RLHF). Modern systems are pipelines of types, not single cells of this table.

AI today — capabilities and honest limitations

Genuinely strong:

- Pattern recognition
- Translation; code generation
- Creative drafting
- Game playing
- Image and video generation

AI today — capabilities and honest limitations

Genuinely strong:

- Pattern recognition
- Translation; code generation
- Creative drafting
- Game playing
- Image and video generation

Still genuinely hard:

- Common-sense and causal reasoning
- Long-horizon planning
- Robustness; calibrated uncertainty
- Knowing *what it doesn't know*

AI today — capabilities and honest limitations

Genuinely strong:

- Pattern recognition
- Translation; code generation
- Creative drafting
- Game playing
- Image and video generation

Still genuinely hard:

- Common-sense and causal reasoning
- Long-horizon planning
- Robustness; calibrated uncertainty
- Knowing *what it doesn't know*

The right column is kind of what our a syllabus is: generalisation theory, energy-based modelling, validation, Bayesian uncertainty... we will try to touch each one.

Future directions and challenges

Technical:

- Artificial General Intelligence?
- Explainable AI
- Data- and energy-efficient learning
- Multimodal + embodied understanding

Societal:

- Jobs and skill displacement
- Bias and fairness
- Privacy and surveillance
- Misinformation at scale

AI for Engineering & Science — where I intends to live

The moral of 70 years — and the case for this course

What every winter killed:

- Products, companies, funding
- Specific architectures and toolchains
- Confident predictions

The moral of 70 years — and the case for this course

What every winter killed:

- Products, companies, funding
- Specific architectures and toolchains
- Confident predictions

What no winter ever killed:

- Linear algebra
- Probability
- The theory of learning and generalisation
- Optimisation

The moral of 70 years — and the case for this course

What every winter killed:

- Products, companies, funding
- Specific architectures and toolchains
- Confident predictions

What no winter ever killed:

- Linear algebra
- Probability
- The theory of learning and generalisation
- Optimisation

Hence: this course

Frameworks will change possibly before you graduate. The mathematics of this course is the part of the field that has survived every boom and every bust since 1956. That is precisely why we teach *fundamentals*.

The moral of 70 years — and the case for this course

What every winter killed:

- Products, companies, funding
- Specific architectures and toolchains
- Confident predictions

What no winter ever killed:

- Linear algebra
- Probability
- The theory of learning and generalisation
- Optimisation

Hence: this course

Frameworks will change possibly before you graduate. The mathematics of this course is the part of the field that has survived every boom and every bust since 1956. That is precisely why we teach *fundamentals*.

The real question

Can someone from *this* class push the frontier of AI forward?

Some of you might already be closer than you think.

Questions!

Next lecture: the Linear Algebra review begins — bring the course map.