



AI Foundations — Lesson 2 Notes

History of Artificial Intelligence | APEP Academy

Learning Objectives

- Trace major stages in the development of artificial intelligence from early ideas about machine intelligence to modern generative AI.
- Explain why Alan Turing, the 1956 Dartmouth workshop, expert systems, machine learning and deep learning are important milestones.
- Recognise the pattern of AI optimism, limitations and renewed investment that contributed to the AI winters.
- Connect AI history to practical decisions entrepreneurs and professionals make when adopting AI tools.

1. What existed before modern AI?

The idea of creating machines that can perform tasks associated with human reasoning is much older than today's chatbots. Long before modern computers, philosophers and inventors explored formal logic, mechanical calculation and the possibility that reasoning could be represented by rules. The development of programmable electronic computers in the 20th century made it possible to test these ideas computationally.

2. Alan Turing and machine intelligence

In 1950, British mathematician Alan Turing published “Computing Machinery and Intelligence.” He reframed the question of whether machines can think into a practical test involving a machine's ability to produce responses that could be difficult to distinguish from those of a human. The work helped establish machine intelligence as a serious field of inquiry and influenced later discussions about evaluation and human-machine interaction.

3. The 1956 Dartmouth workshop

The Dartmouth Summer Research Project on Artificial Intelligence, held in 1956, is widely associated with the formal emergence of AI as a research field. Researchers proposed that aspects of learning and intelligence could be described precisely enough for machines to simulate them. The term “artificial intelligence” became closely associated with this research programme.

4. Early symbolic and rule-based AI

Early AI systems often represented knowledge explicitly through symbols, rules and logical procedures. Instead of learning patterns from large datasets, a system might be programmed with instructions such as “if condition A is true, perform action B.” This approach worked well for narrowly defined problems but became difficult to scale

when real-world situations contained ambiguity, exceptions and enormous numbers of possible states.

5. Expert systems

During the 1970s and 1980s, expert systems became an important commercial and research application of symbolic AI. They attempted to reproduce parts of a specialist's decision process by combining a knowledge base with an inference engine. They demonstrated that AI could create business value in constrained domains, while also revealing the cost of maintaining large rule sets.

6. AI winters

AI history includes periods commonly called “AI winters,” when expectations exceeded technical progress and funding or public enthusiasm declined. Two widely discussed periods occurred from the mid-1970s into the early 1980s and from the late 1980s into the 1990s, although the exact boundaries vary by source. A useful business lesson is that emerging technology can be valuable without every prediction about it being correct.

7. The rise of machine learning

Machine learning shifted part of the AI problem from manually specifying every rule toward allowing systems to learn patterns from data. Statistical methods, better datasets and increasing computing power made this approach increasingly practical. Instead of telling a system every rule for recognising spam, for example, engineers could train a model on examples of spam and legitimate messages.

8. The deep-learning revolution

Deep learning uses multi-layer neural networks to learn increasingly complex representations from data. The approach benefited from larger datasets, graphics processors and improved training techniques. A major public milestone was the strong performance of the AlexNet system in the 2012 ImageNet competition, which helped accelerate interest in deep neural networks for computer vision.

9. Transformers and modern generative AI

The 2017 research paper “Attention Is All You Need” introduced the Transformer architecture. Transformers became foundational to many modern language and multimodal AI systems because they can model relationships among elements in sequences efficiently. Generative AI systems can produce text, images, audio, code and other media from learned patterns. The public release of ChatGPT in November 2022 helped bring conversational generative AI into mainstream use.

10. AI in everyday life

AI is now embedded in many ordinary workflows: search ranking, recommendations, fraud detection, translation, navigation, customer support, spam filtering, image processing and productivity software. This matters because AI is not limited to a single industry. The practical question is increasingly how to identify useful tasks, choose appropriate tools and introduce them responsibly.

11. Why AI history matters to entrepreneurs and professionals

- Technology hype is not the same as business value. Look for measurable problems that AI can solve.
- The best tool depends on the task, data, workflow, risk and user—not simply on which model is newest.

- AI capabilities change quickly, so professionals need a habit of testing, evaluating and updating their workflows.
- Human judgement remains important for quality control, context, ethics, privacy and consequential decisions.
- Successful adoption is usually a workflow problem as much as a technology problem.

Key Takeaways

- AI developed through several overlapping approaches rather than one straight line of progress.
- Symbolic AI focused heavily on rules and explicit knowledge; machine learning shifted more emphasis toward learning from data.
- Deep learning expanded what neural networks could accomplish when combined with large datasets and powerful computing.
- Transformers became a major foundation for modern language and generative AI.
- Understanding AI's history helps you separate durable principles from short-lived hype.

AI Prompt Practice

Prompt: “Explain the history of AI to a small-business owner who has never studied computer science. Use a simple timeline, explain five major turning points, and finish with three lessons the owner can apply when choosing AI tools.”

Challenge: Improve the prompt by specifying the audience, desired format, level of detail, examples and a requirement to distinguish historical facts from interpretation.

Lesson Summary

Artificial intelligence has developed through cycles of ambitious ideas, technical constraints, new methods and improved computing resources. The journey from symbolic systems to machine learning, deep learning and generative AI shows why AI adoption should be approached with both curiosity and discipline. For APEP learners, the practical lesson is to understand the technology well enough to make better business and workflow decisions.

Knowledge Check

1. Why is the 1956 Dartmouth workshop important in AI history?
2. What is the basic idea behind symbolic or rule-based AI?
3. What problem did expert systems attempt to solve?
4. What is meant by an “AI winter”?
5. How does machine learning differ from manually writing every rule?
6. Why was the 2012 AlexNet result significant?
7. What is a Transformer in the context of modern AI?
8. Give two examples of AI used in everyday life.
9. What is one business lesson that can be learned from AI's history?
10. Improve the Lesson 2 practice prompt by adding one useful constraint.

Lesson Resources

Use the accompanying Notes, AI Prompt Practice Worksheet and Quick Reference Guide for revision and practice.