



Rise of the Machines: The Emergence of Artificial Intelligence

An Online Continuing Education Course for Engineers

Course Number: ET-2080

Credit: 2 Hours / 2 PDH / 2 CPD

Rise of the Machines: The Emergence of Artificial Intelligence

Mark A. Strain, P.E.

Table of Contents

Introduction	4
History.....	4
Alan Turing and the Birth of the Computer Age (1936–1950)	4
Early Sparks: The Dawn of AI (1950s)	4
The AI Boom and the Dartmouth Conference (1956–1970s).....	5
The Rise of Expert Systems and Knowledge-Based AI (1970s–1980s)	5
AI Winters: Hype Meets Harsh Reality (1970s–1990s).....	5
The Machine Learning Resurgence (1990s–2010s).....	5
AI in the Real World: Mainstream Adoption (2010s–2020s).....	6
The Age of Generative AI and Ethical Reckoning (2020s–Present)	6
The Turing Test	6
Concept and Setup: The Imitation Game	6
Purpose and Philosophical Significance.....	7
Examples of Questions in a Turing Test.....	7
Criticism and Limitations	7
Real-World Milestones and Attempts	8
The Turing Test in the Modern AI Era.....	8
More Than a Test.....	8
Does ChatGPT Pass the Turing Test (according to ChatGPT).....	9
Strengths of ChatGPT in Relation to the Turing Test	9
Limitations.....	9
Conclusion	9
Moore’s Law.....	9
Asimov’s Three Laws of Robotics.....	10
ChatGPT	11
What is ChatGPT (according to ChatGPT).....	11
How ChatGPT “Thinks” — The Human Analogy.....	12
The Generative Part	12
The Pre-trained Part.....	12
The Transformer Part	12
ChatGPT Uses.....	13

Technological Singularity	14
Fermi’s Paradox	14
The Paperclip Apocalypse	15
How It Happens:	15
Lesson:	15
Questions for ChatGPT.....	15
What do you think is the near-term evolution of chatbots?.....	15
Do you think chatbots will ever become harmful?	16
Will chatbots ever be made so that they can improve themselves without human intervention?.....	16
Do you think that chatbots could ever become harmful or destructive?	17
Will chatbots ever become sentient?.....	17
If a chatbot became sentient, what could we expect?	17
Are you sentient?	18
Is your neural network based on neurons with a sigmoid function?.....	18
News Articles on the Dangers of AI	18
"Bing's AI bot tells reporter it wants to 'be alive', 'steal nuclear codes' and create 'deadly virus'"	18
"AI system resorts to blackmail when its developers try to replace it"	20
Summary	21
References	22

Introduction

Humanity is on the cusp of unthinkable changes. Technology has the potential to rapidly progress in the next few decades in ways that we cannot even imagine. Artificial intelligence (AI) has already emerged in the form of chatbots whose conversations are indistinguishable from a human response.

Bots have been around for years, slowly creeping into websites in the help menus and into voice prompts posing as human messengers to answer our questions and guide us to pre-canned answers. Most of us have experienced the frustration of navigating these artificial messengers hurrying through the menu machine to finally get to a real human. You could tell right away that the voice on the other side of the conversation was synthetic, but now it is getting increasingly difficult to differentiate what is human and what is machine. Bots are computer programs that simulate human conversation through voice commands or text chats. A bot is an autonomous program that can interact with systems or people.

Recent developments in machine learning algorithms, such as deep learning and neural networks, perform AI tasks such as natural language generation, image recognition, speech recognition, and text-to-speech synthesis. We are approaching the technological singularity, where technological advancements accelerate so rapidly that they could surpass human comprehension and control. As with other unknown ventures, this is a scary prospect.

History

Alan Turing and the Birth of the Computer Age (1936–1950)

In 1936, British mathematician and logician Alan Turing introduced the groundbreaking concept of a "universal machine," now known as the Turing Machine. This theoretical construct could simulate any algorithmic process, laying the foundational principles of modern computer science and computability theory. It wasn't just a conceptual tool—it provided a framework for thinking about the limits of mechanical computation, foreshadowing the programmable digital computers we rely on today.

In 1950, Turing posed a provocative question in his seminal paper "Computing Machinery and Intelligence: Can machines think?" He proposed the Imitation Game, now known as the Turing Test, as a way to assess machine intelligence. If a machine could engage in a text-based conversation indistinguishable from a human, it might be said to exhibit intelligent behavior. This idea would become a philosophical and technical cornerstone of artificial intelligence.

Early Sparks: The Dawn of AI (1950s)

The 1950s saw the first practical attempts to build intelligent machines. Researchers developed early AI programs capable of solving algebraic equations, playing simplified games like chess, and proving logical theorems. Notably, Allen Newell, Herbert A. Simon, and Cliff Shaw created the Logic Theorist (1955), often considered the first AI program. These systems operated on symbolic logic—using rules to manipulate abstract symbols much like a human reasoner.

The AI Boom and the Dartmouth Conference (1956–1970s)

The field of artificial intelligence was officially born at the Dartmouth Summer Research Project on Artificial Intelligence in 1956, organized by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon. Here, McCarthy coined the term “artificial intelligence,” launching a new era of computer science focused on replicating human-like reasoning.

Throughout the 1960s and early 1970s, optimism surged. AI programs like the General Problem Solver could tackle puzzles and logical tasks, and rudimentary natural language systems such as ELIZA demonstrated the illusion of conversational understanding. These early systems operated using symbolic AI, a top-down approach that relied heavily on handcrafted rules and explicit knowledge representation.

The Rise of Expert Systems and Knowledge-Based AI (1970s–1980s)

By the 1970s, attention turned to expert systems—software designed to emulate the decision-making abilities of human specialists. Systems like MYCIN (for medical diagnosis) and DENDRAL (for chemical analysis) showcased how rule-based logic could outperform non-expert humans in narrow domains. These applications ignited corporate and government interest, as AI seemed poised to revolutionize professional industries.

AI Winters: Hype Meets Harsh Reality (1970s–1990s)

Despite early promise, the limitations of symbolic AI became apparent. Systems struggled with ambiguity, learning from experience, and applying knowledge outside predefined scenarios. As expectations outpaced progress, the field suffered multiple setbacks known as AI Winters—periods marked by declining funding, waning public interest, and academic disillusionment.

One contributing factor was the stagnation of neural network research. Although the perceptron, an early neural model, had shown potential, researchers could not effectively train multi-layer networks, leading to skepticism about the viability of connectionist approaches.

The Machine Learning Resurgence (1990s–2010s)

The 1990s marked a turning point. Researchers shifted focus from rule-based AI to machine learning, where computers learn from data rather than being explicitly programmed. The development of algorithms such as support vector machines, decision trees, and Bayesian networks allowed for more flexible, data-driven models.

A breakthrough came with the revival of neural networks and the emergence of deep learning in the 2010s. Fueled by exponential increases in computing power (especially GPUs), access to vast datasets, and new techniques like backpropagation, deep neural networks achieved stunning results in image classification (e.g., ImageNet), speech recognition, and natural language processing. Landmark achievements like AlphaGo’s victory over a human Go champion in 2016 captured the public imagination and proved that AI could master complex, intuition-based tasks.

AI in the Real World: Mainstream Adoption (2010s–2020s)

AI became woven into daily life. Virtual assistants like Siri, Alexa, and Google Assistant interpreted voice commands; self-driving cars navigated roads using AI-powered vision and planning systems; and recommendation algorithms curated music, videos, and shopping choices. AI applications expanded across healthcare, finance, entertainment, and manufacturing, becoming an invisible but vital part of modern infrastructure.

The Age of Generative AI and Ethical Reckoning (2020s–Present)

The 2020s ushered in a new frontier: generative AI. Large language models (LLMs) like ChatGPT, GPT-4, and Claude demonstrated an uncanny ability to generate human-like text, answer questions, write code, and engage in conversation. Visual models like DALL-E, Midjourney, and Stable Diffusion generated art and images from text prompts, opening new creative and commercial possibilities.

However, this power brought new challenges. AI's ability to generate convincing but incorrect or biased content, contribute to misinformation, and displace jobs triggered intense debates about AI ethics, transparency, and accountability. Scholars and policymakers began grappling with questions of safety, fairness, and social impact.

In response, international bodies and governments are now drafting AI regulations, frameworks, and ethical guidelines. The goal is to ensure that AI development aligns with human values and serves the public good.

The Turing Test

The Turing Test, proposed in 1950 by the brilliant British mathematician and computer scientist Alan Turing, remains one of the most iconic and thought-provoking ideas in the history of artificial intelligence. Introduced in his seminal paper "Computing Machinery and Intelligence," the test was conceived as a practical way to approach the age-old philosophical question: "Can machines think?" Rather than getting entangled in abstract definitions of "thought" or "consciousness," Turing suggested evaluating machine intelligence by observing whether its behavior is indistinguishable from that of a human.

Concept and Setup: The Imitation Game

Turing framed his proposal as a type of Imitation Game, which involves three participants:

1. **A human interrogator (or judge)**, tasked with determining which of the two respondents is human.
2. **A human respondent**, who answers questions honestly and naturally.
3. **A machine (or AI system)**, which attempts to imitate a human convincingly enough to fool the interrogator.

To ensure objectivity, the communication occurs through text-only exchanges, stripping away any physical cues like voice, facial expressions, or appearance. The interrogator poses a series of questions—

ranging from logical puzzles to emotionally charged topics or everyday observations—and evaluates the responses.

The goal of the machine is not to give correct answers, but to give human-like ones. If the interrogator cannot reliably tell which respondent is the machine based on their answers, the machine is considered to have passed the Turing Test. This doesn't mean the machine is intelligent in a human sense—it means its behavior is functionally indistinguishable from that of a person.

Purpose and Philosophical Significance

Turing's genius lay in reframing the problem. Rather than asking whether machines can actually think—a question mired in metaphysics—he asked whether machines can act as if they think. This operational definition shifted the focus from internal mechanisms to external behavior.

In doing so, Turing laid the foundation for modern artificial intelligence: intelligence as judged by observable actions.

Examples of Questions

To evaluate the machine's responses, interrogators ask questions designed to test various dimensions of human-like behavior.

- "What is your opinion on the current state of the world?" (tests self-awareness, and reasoning.)
- "Can you describe a complex task in simple terms?" (tests communication, and plausibility.)
- "The more you know about this problem, the more you should be able to solve it." (tests problem-solving, and logical coherence.)
- "How would you feel about this situation?" (tests emotional depth and empathy.)

To view the remainder of the course material and to take the quiz for PDH credit, you must purchase the course.

Close this window and click "Add to cart" on the product page.

The machine must not only provide linguistically correct answers but also display contextual awareness, emotional nuance, and logical coherence.

Criticism and Limitations

Although revolutionary, the Turing Test is far from flawless. Over the decades, many scholars have raised valid criticisms:

- Shallow Imitation vs. True Understanding: A machine can pass the test using clever scripting, pre-programmed phrases, or probabilistic models without truly understanding what it is saying.
- Mimicry is not Intelligence: Deception can mask emptiness. Human-like responses don't imply self-awareness, insight, or intentionality.
- Narrow Scope: The test focuses solely on language and excludes other forms of intelligence, such as visual perception, creativity, physical intuition, or mathematical reasoning.