

AI: Our Quest to Understand the Nature of Intelligence Itself

Code Smarter, Not Harder: Unlocking AI-Assisted Coding
Coffee, Cookies, and Coding (C-cubed) Workshops

March 23, 2026


“The story of AI is an exhilarating saga of new programs and new hardware, yet it is also the story of a slow but steady acquisition of knowledge about how humans think.” – Daniel Crevier ¹

The concept of artificially begetting life and intelligence has ancient roots. Throughout most of human history, such creation was attributed to divine power or alchemical mysticism ²⁻⁴. The Enlightenment and First Industrial Revolution marked a pivotal shift: an increasingly secularized society posited that the human body and mind were scientifically describable and replicable ^{5,6}. This transition laid the conceptual foundation for contemporary efforts to develop intelligence artificially through machines.

Our pursuit of artificial intelligence reveals fundamental human desires (perfection, relief from rote work), fears (creations rebelling against creators), and spurs moral and societal questions (what would happen if machines become indistinguishable from humans?). Humanity has grappled with these themes and philosophical questions long before contemporary AI discourse emerged.

For most of history, we lacked the technology and science to attempt creating artificial intelligence. Today, our limitations stem less from technical constraints than from incomplete understanding of the origins of consciousness, creativity, and thought itself.

Amidst the humdrum and hype lies a story that reveals much about ourselves and a transformative, powerful technology in its own right, even though it's far from achieving true artificial intelligence. Let's meet the real AI so you can ride the wave of excitement without getting pulled under.

 Discussion Question

What did Rosenblatt get right and wrong, how did his perspective shape AI, and do today's prominent figures still think this way?



Figure 1: New York Times "Times Machine" article from Tuesday July 8th, 1958 page 25.

A Field Guide to AI Typologies and Systems

All existing AI is weak AI (Type I or II), while “strong” and “super” AI remain theoretical with no field consensus on their achievability. Type II systems currently dominate the market due to their flexibility, natural user interaction, and ability to adapt based on recent input ^{4,7}.

Notably, I intentionally qualify or avoid the terms “learning” or “understanding” when describing “weak” AI, as it overstates these systems’ capabilities. Though commonly used in the field, it is more accurately a mathematical process of adjusting weights and parameters, also referred to as “tuning,” “optimizing,” or “updating.”

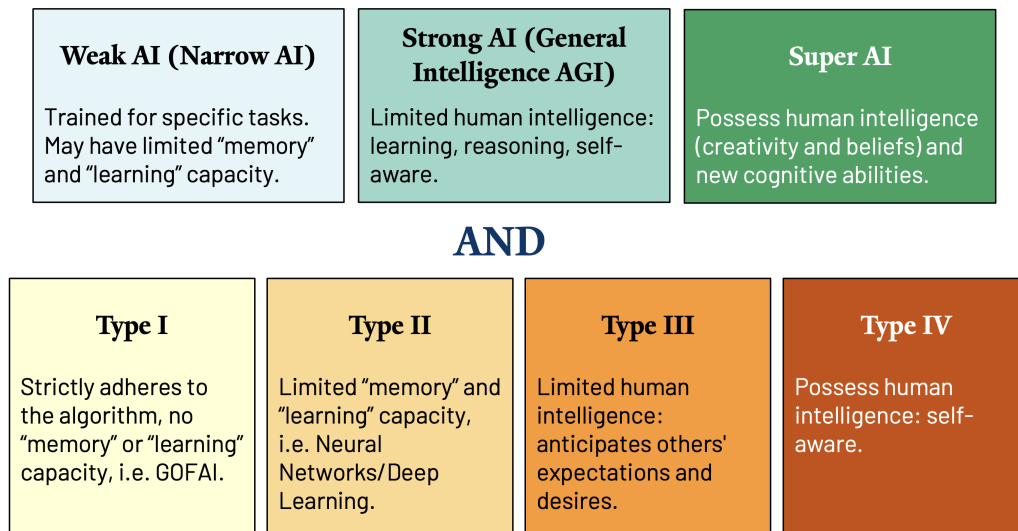



Figure 2: AI categories and types with definitions.

 Discussion Question

Identify the applicable category and type for each AI system below (both theoretical and actual). Note that “Expert Systems” refers to “Good Old-Fashioned AI (GOFAL),” which saw a brief resurgence during the 20th century AI winters.

<p>Expert Systems</p> <hr style="width: 50%; margin: 0 auto;"/>	<p>Star Wars’ C- 3PO</p> <hr style="width: 50%; margin: 0 auto;"/>	<p>ChatGPT/Claude/etc.</p> <hr style="width: 50%; margin: 0 auto;"/>
---	--	--

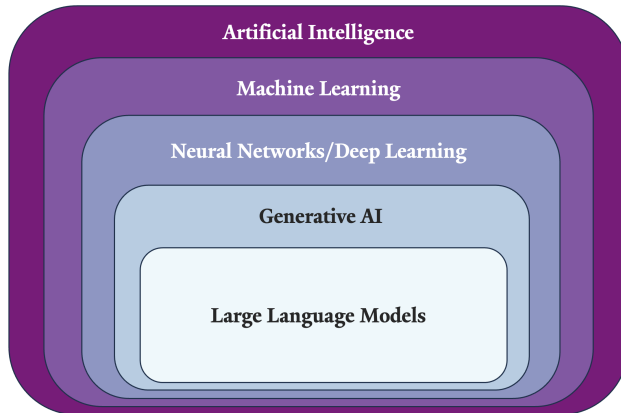


Figure 3: Venn diagram illustrating different AI systems. Adapted from Figure 1-3 (page 5) in *AI-Assisted Programming - Better Planning, Coding, Testing, And Deployment* by Tom Taulli.

Artificial intelligence is the broadest category, encompassing all AI systems (theoretical and actual), including “weak,” “strong,” and “super” classifications, as well as Types I-IV⁷.

Machine learning is a category of Type II “weak” AI that uses mathematical techniques to optimize predictions modeled on training data, rather than relying on explicitly coded logic. Some systems can “learn” by iteratively adjusting model parameters to reduce prediction errors when evaluated against new or unseen data^{7,8}.

Training approaches can take different forms, yet most stem from or are related to three core methods:

- Supervised learning: using labeled data showing correct answers.
- Unsupervised learning: detecting patterns through clustering and other techniques.
- Reinforcement learning: trial and error via environmental rewards and penalties, such with sensor on self-driving cars.

Neural networks are a subset of machine learning that use pattern recognition techniques loosely inspired by biological neuron propagation. Contemporary networks typically leverage many layers – giving rise to the term **deep learning**^{4,7,9,10}. Due to their abstract mathematics and probabilistic nature, their decision-making process is difficult to interpret or explain in real-world terms, which is why AI is often referred to as a “black box.”

Generative AI refers to neural networks that create new content (images, audio, video, text) based on learned patterns. This advancement occurred after the development of **Natural Language Processing (NLP)**, a machine learning-based modeling technique that trains computers to generate responses relevant to a given prompt in a human-esque manner¹¹. These models train on massive datasets using unsupervised deep learning. They often undergo fine-tuning through reinforcement learning from human feedback (RLHF) to shift from generalist to specialized performance with higher accuracy and fidelity^{7,12}.

Large Language Models (LLMs) are generative AI models specialized for text generation. Popular foundation models include GPT, Gemini, Claude, and LLaMA—tools we’ll use for code

generation through chatbots or IDE integrations^{7,13}.

Neural networks are loosely inspired by biological neurons, with mathematical units called perceptrons serving as their counterpart (Figure: 4)¹⁴. Between each layer, signals are modified by weights (scaling importance) and biases (shifting importance), then passed through non-linear activation functions that determine what gets propagated to the next layer. This continues until a final activation function on the output layer generates the prediction \hat{Y} ^{9,10,15}.




Figure 4: Simplified diagram of a neural network. Were more hidden layers shown, this would make it a deep learning neural network diagram. Adapted from the “Feedforward vs Backpropagation ANN” LinkedIn post by SaffronEdge.

The feedforward pass first processes signals from input to output, after which backpropagation works in reverse – efficiently identifying where weights and biases




Figure 5: Modified version of Figure 5.1 (page 228) in *Pattern Recognition and Machine Learning* by Christopher Bishop.

should be adjusted to minimize prediction error^{9,15}. Backpropagation uses the results of the feedforward process to optimize the model, generating gradient descents to find possible optimal solutions. Unlike some mathematical setups that guarantee a single optimal solution, this is not the case with neural networks^{10,16}.

I’ll illustrate the feedforward process with a simplified example from Bishop, leaving backpropagation for you to explore

on your own¹⁰. Consider the neural network in Figure 5, which has one hidden layer and two transformation layers¹⁰. Reading from the input to output layer $i = 1, \dots, D$, $j = 1, \dots, M$, and $k = 1, \dots, K$ represent the unique perceptrons in each layer. The signal from the input layer (x_1, \dots, x_D) gets weighted $(w_{ji}^{(1)})$ and biased $(w_{j0}^{(1)})$ to create a_j :

$$a_j = \sum_{i=1}^D w_{ji}^{(1)} x_i + w_{j0}^{(1)} \quad (5.2)$$

Then the layer-to-layer, nonlinear activation function $h(\cdot)$ (i.e. logistic sigmoid or hyperbolic tangential) is applied:

$$z_j = h(a_j) \quad (5.3)$$

Using the first layer transformation's output, z_j , we again apply weights and biases linearly:


$$a_k = \sum_{j=1}^M w_{ki}^{(2)} z_j + w_{k0}^{(2)} \quad (5.4)$$

A new activation function, σ , suited to the nature of the data is used to translate the last hidden layer's output, a_k , to the final output, y_k :


$$y_k = \sigma(a_k) \quad (5.5)$$

The overall network function for the feedforward process of this neural network yields:

$$y_k(\mathbf{x}, \mathbf{w}) = \sigma \left(\sum_{j=1}^M w_{kj}^{(2)} h \left(\sum_{i=1}^D w_{ji}^{(1)} x_i + w_{j0}^{(1)} \right) + w_{k0}^{(2)} \right) \quad (5.7)$$

 Discussion Question

Do we agree with the definitions for AI and Machine Learning recorded in The Code of Laws of the United States of America?



15 U.S. Code § 9401. Definitions Date : January 1, 2021

(3) Artificial intelligence
 The term "artificial intelligence" means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to-

- (A) perceive real and virtual environments;
- (B) abstract such perceptions into models through analysis in an automated manner;

and

- (C) use model inference to formulate options for information or action.

(11) Machine learning
 The term "machine learning" means an application of artificial intelligence that is characterized by providing systems the ability to automatically learn and improve on the basis of data or experience, without being explicitly programmed.

10 U.S. Code § 2358 - Renumbered § 4001 Date : January 13, 2021

(g) Artificial Intelligence Defined

1. Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.
2. An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
3. An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
4. A set of techniques, including machine learning, that is designed to approximate a cognitive task.
5. An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

Figure 6: U.S. code definitions codified by the House of Representatives for AI and related mathematical techniques: [10 USC 2358](#) and [15 USC 9401](#).

Glossary

Term	Definition
Artificial Intelligence (AI)	Intelligence or mind displayed by artificial life or machines, analogous to the natural intelligence of animals and humans, capable of perceiving its environment and taking action. AI can be classified by category ("weak," "strong," or "super") and by type (Types I-IV), representing varying degrees to which the technology has achieved true artificial intelligence.
Automaton, Automata	A self-moving mechanical device, typically resembling an animal or human, that operates without direct human control. Some automata follow predetermined instructions, while others can produce a range of responses to varying circumstances. Might be considered an early Type I "weak" AI.
Black Box	Complex entity, machine, or system whose outputs are known or observable but whose inner contents and internal workings are hidden, unknown, opaque, and mysterious to the user.
Big Data	A shift in data collection marked by increased volume and velocity (speed of collection), where data expanded beyond traditional tabular formats to include greater variety: images, free text, audio, and spatial data. Maintaining data veracity became more challenging, increasing the need for valuable (i.e. reliable) data. Big Data also exacerbates the "curse of dimensionality," where the number of variables (columns) outpaces the number of observations (rows).
Central Processing Unit (CPU)	Electronic circuitry that run a computer's operating system and apps and manage a variety of other computer operations. Thought of as the "brain" of the computer.
Connectionism	A movement in cognitive science that hopes to explain intellectual abilities using artificial neural networks (also known as "neural networks" or "neural nets").
Computer Vision	A field of AI that enables machines to interpret and understand visual information from the world, such as images and video.

Term	Definition
Deep Networks	A type of neural network characterized by many hidden layers. Made computationally feasible through theoretical advances such as backpropagation, deep networks support deep learning-based machine learning – a core technique underlying generative AI.
Disambiguate	To resolve ambiguity by clarifying or distinguishing the intended meaning of a word, phrase, or concept. “Weak” AI algorithms lack the ability to understand context beyond modeled patterns and typically underperform in this regard.
Generative AI	A type of AI based on neural network/deep learning machine learning that generates new content, such as text, images, audio, or video, by learning patterns from existing data.
Good Old-Fashioned AI (GOF AI)	An early approach to AI based on explicit, hand-coded rules and logical reasoning. All Type I “weak” AI are GOF AI.
Graphics Processing Units (GPU)	A computer circuit designed for fast graphic and image processing. Its greater number of processing cores and memory also make it ideal for complex mathematical calculations and were necessary in the advancement of AI development.
Large Language Models (LLMs)	Generative AI models designed specifically for text generation, i.e. chat bots. Popular foundation models behind these tools include ChatGPT, Gemini, Claude, and LLaMA.
Machine	A mechanical structure or device based on one or more components (such as lever, pulley, wheel and axle, inclined plane, screw, wedge) that changes the direction or magnitude of a force.
Machine Learning	“Weak” AI with Type II capabilities; limited memory and the ability to optimize based on new input. Canonically refers to mathematical techniques where an algorithm optimizes criteria to minimize prediction errors based on answers it generates from training data.
Mechanism, Mechanical	Something made of parts that move or work together to perform an action; a machine or something resembling a machine.

Term	Definition
Natural Language Processing (NLPs)	A technology that uses machine learning-based modeling to train computers to generate responses relevant to a given prompt in a human-esque manner, effectively emulating human communication. This advancement served as a foundational development that enabled the field of generative AI.
Neural Networks	A subset of machine learning that use pattern recognition techniques loosely inspired by biological neuron propagation. Contemporary networks typically leverage many layers – giving rise to the term deep learning . Due to their abstract mathematics and probabilistic nature, their decision-making process is difficult to interpret or explain in real-world terms, which is why AI is often referred to as a “black box.”
Perceptron	An alternative term used to describe individual “neurons,” the most basic unit of a neural network, modeled loosely after a biological neuron in the human brain. It is used to distinguish the differences between machine learning-based neural networks and biological neural connections.
Programmed	Supplied with a predetermined set of (coded) instructions for automatic performance.
Reinforcement learning with human feedback (RLHF)	Human users respond to generated content with evaluations and scores the model can use to update the model for greater accuracy or relevance.
Robot, Bot	Complex and ambiguous to define, but a robot usually is a machine or self-moving object with a power source that provides energy. It can be “programmed” to “sense” its surroundings, and has a kind of “intelligence” or way of processing data to “decide” to interact with the environment to perform actions or tasks. Talos, the bronze animated statue powered by ichor, fits this definition.
Strong AI	A theoretical AI capable of learning, understanding, and reasoning at a human level.
Super AI	A theoretical AI capable of creativity and holding belief systems at a human level. Also, might possess additional cognitive abilities that surpass human capacity.

Term	Definition
Weak AI	Also referred to as “narrow AI,” a system designed to perform specific tasks, either explicitly defined or determined by probabilistic models from detected patterns in a training dataset. Though some systems can mimic human-esque responses, they lack genuine understanding or reasoning capacities. Type II “weak” AI has limited “memory” and “learning” capacity –that is, the ability to tune model parameters based on recent input.

Appendix

This document was created using Quarto, which uses Pandoc as a primary method of parsing BibTeX citations. Unfortunately, this does not compile with LaTeX formatted citations. Where LaTeX was used (such as formatting pictures and tables), references were either embedded in nearby, relevant text or in this appendix.

New York Times references: ¹⁷

U.S. Code references: ¹⁸⁻²²

Glossary references: ^{4,7-13,23-30}

References

1. Crevier, D. *AI: The Tumultuous History of the Search for Artificial Intelligence*. (Basic Books, New York, NY, 1993).
2. Alex Shashkevich. Stanford researcher examines earliest concepts of artificial intelligence, robots in ancient myths. Stanford report - arts & humanities. <https://news.stanford.edu/stories/2019/02/ancient-myths-reveal-early-fantasies-artificial-life> (2019).
3. E.R. Truitt. Surveillance, companionship, and entertainment: The ancient history of intelligent machines. The MIT press reader. <https://thereader.mitpress.mit.edu/the-ancient-history-of-intelligent-machines/> (2021).

4. Mayor, A. *Gods and Robots: Myths, Machines, and Ancient Dreams of Technology*. (Princeton University Press, Princeton, 2018).
5. Hatfield, G. René descartes. in *Stanford encyclopedia of philosophy archive* (Stanford, 2014).
6. Kantosalo, A., Falk, M. & Jordanous, A. Embodiment in 18th century depictions of human-machine co-creativity. *Front Robot AI* **8**, 662036 (2021).
7. Taulli, T. *Artificial Intelligence Basics: A Non-Technical Introduction*. (Apress, Berkeley, CA, 2019). doi:10.1007/978-1-4842-5028-0.
8. Dave Bergmann. What is machine learning? IBM think topics. <https://www.ibm.com/think/topics/machine-learning> (2025).
9. Fangfang Lee. What is a neural network? IBM think topics. <https://www.ibm.com/think/topics/neural-networks> (2021).
10. Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. (Springer, 2006).
11. Cole Stryker & Jim Holdsworth. What is NLP (natural language processing)? IBM think topics. <https://www.ibm.com/think/topics/natural-language-processing> (2024).
12. Cole Stryker & Mark Scapicchio. What is generative AI? IBM think topics. <https://www.ibm.com/think/topics/generative-ai> (2024).
13. Cole Stryker. What are large language models (LLMs)? IBM think topics. <https://www.ibm.com/think/topics/large-language-models> (2021).
14. SaffronEdge. Feedforward vs backpropagation ANN. LinkedIn. <https://www.linkedin.com/pulse/feedforward-vs-backpropagation-ann-saffronedge1/> (2023).
15. English, M. L. in P. Deep learning course – lesson 5: Forward and backward propagation. Medium. <https://medium.com/@nerdjock/deep-learning-course-lesson-5-forward-and-backward-propagation-ec8e4e6a8b92> (2023).
16. Dave Bergmann & Cole Stryker. What is backpropagation? IBM think topics. <https://www.ibm.com/think/topics/backpropagation> (2024).
17. New navy device learns by doing. *The New York Times* 25 (1958).
18. Live streaming of the house. Live streaming of the house. <https://live.house.gov/> (2026).
19. United states house of representatives seal. *Wikipedia* (2026).
20. United States House of Representatives. 10 USC 2358: Research and development projects. US law codebook. <https://uscode.house.gov/view.xhtml?hl=false&edition=2020&req=granuleid/%3AUSC-2021-title10-section2358&num=0> (2021).
21. United States House of Representatives. 15 USC 9401: definitions. US law codebook. [https://uscode.house.gov/view.xhtml?req=\(title:15/%20section:9401/%20edition:prelim\)\(2026\)](https://uscode.house.gov/view.xhtml?req=(title:15/%20section:9401/%20edition:prelim)(2026)).

22. Artificial intelligence (AI). LII / legal information institute. [https://www.law.cornell.edu/wex/artificial_intelligence_\(ai\)](https://www.law.cornell.edu/wex/artificial_intelligence_(ai)).
23. Kalsi, H. S. What is good old-fashioned AI? Medium. <https://medium.com/@harrpreet/what-is-good-old-fashioned-ai-59ff4a4719ce> (2024).
24. Annie Badman & Matthew Kosinski. What is big data? IBM think topics. <https://www.ibm.com/think/topics/big-data> (2024).
25. Mesh Flinders, Stephanie Susnjara & Ian Smalley. What is a GPU? IBM think topics. <https://www.ibm.com/think/topics/gpu> (2025).
26. Tom Taulli. *AI-Assisted Programming - Better Planning, Coding, Testing, and Deployment*. (O'Reilly Media, Sebastopol, CA, 2024).
27. Phill Powell & Ian Smalley. What is a central processing unit (CPU)? IBM think topics. <https://www.ibm.com/think/topics/central-processing-unit> (2024).
28. Buckner, C. & Garson, J. **Connectionism**. in *The stanford encyclopedia of philosophy* (eds. Zalta, E. N. & Nodelman, U.)(Metaphysics Research Lab, Stanford University, 2019).
29. Chess, J. Generative artificial intelligence: Levels of artificial intelligence. U.s. Military academy library. <https://library.westpoint.edu/GenAI/capabilities> (2025).
30. Sospeter Mong'are. Strong AI vs. Weak AI. DEV community. <https://dev.to/msnmongare/strong-ai-vs-weak-ai-1gi1> (2025).