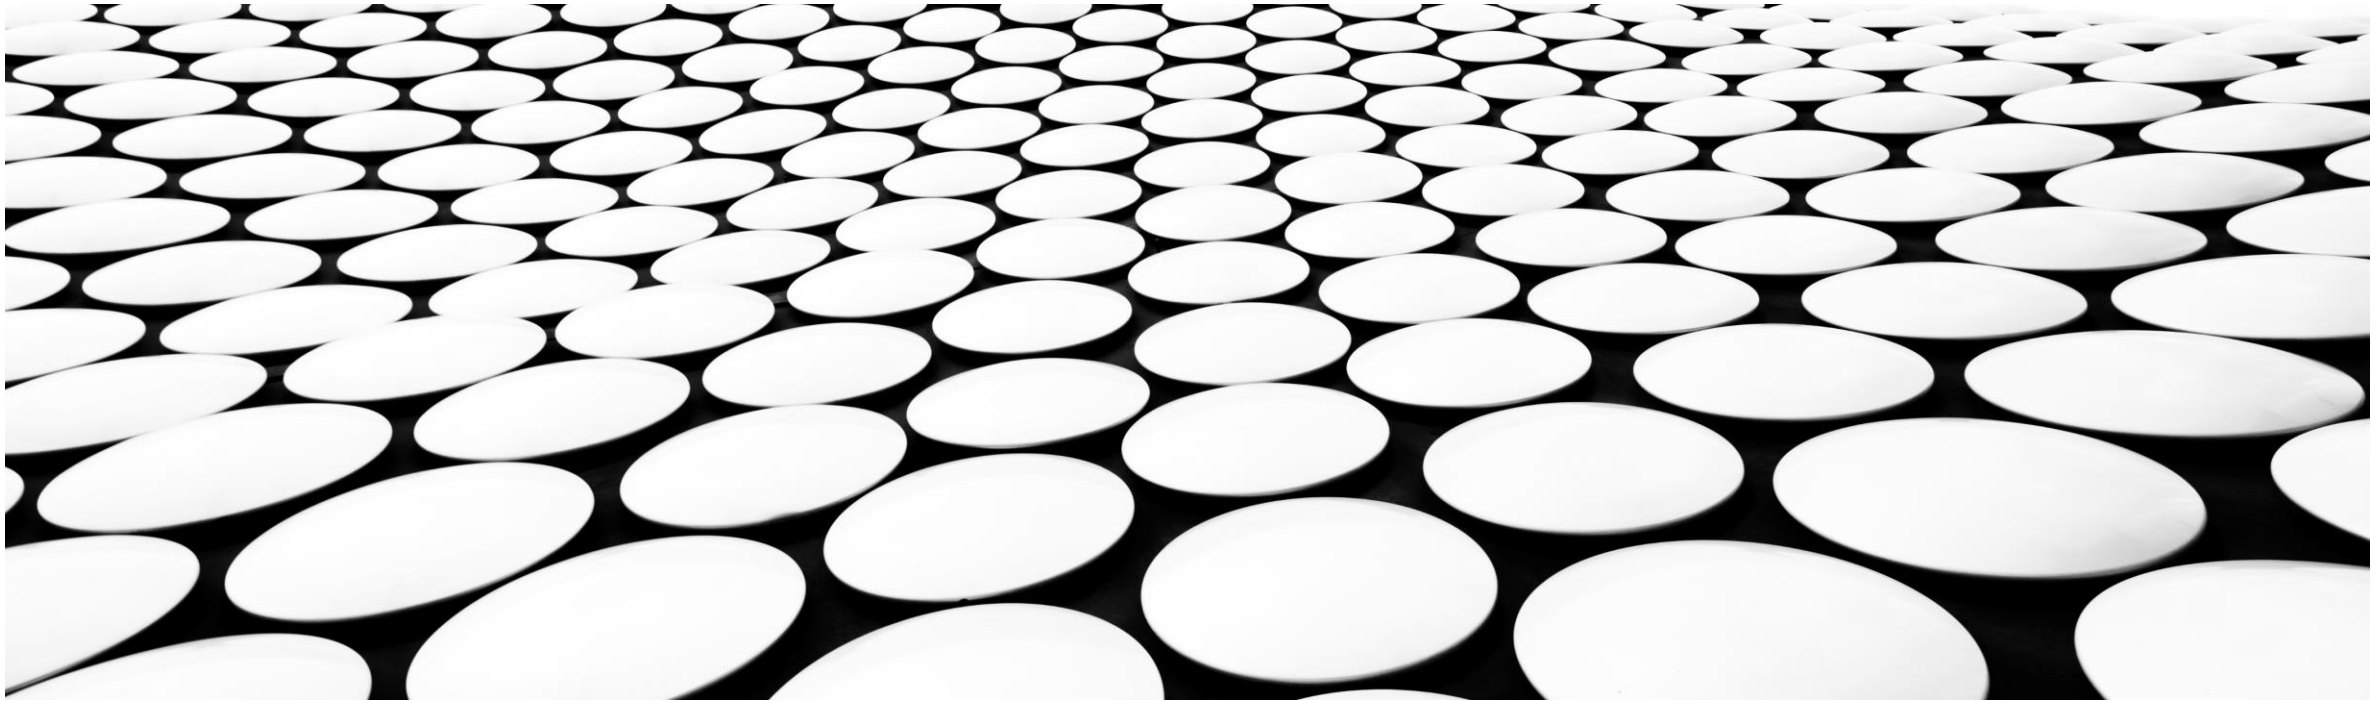
ARTIFICIAL INTELLIGENCE : HISTORY TO FUTURE

JEFF GREENBERG, PHD

NATIONAL DIALOGUES ON BEHAVIORAL HEALTH CONFERENCE. 9 NOVEMBER 2026



PRESENTER

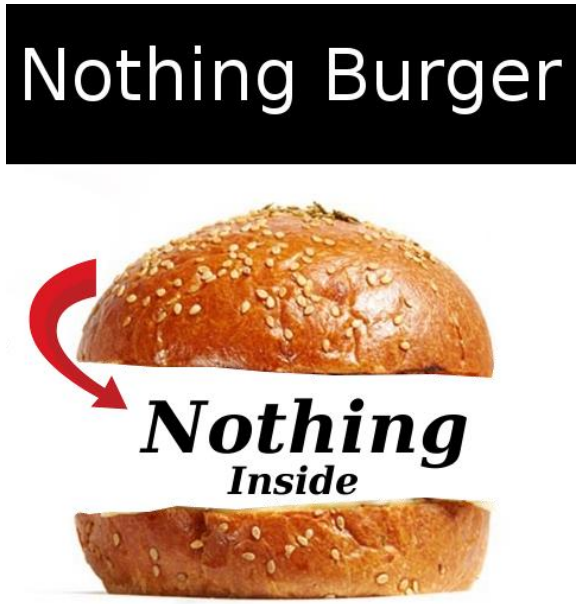
- **Jeff Greenberg, PhD, is a clinical psychologist and senior health executive known for leading innovation at the intersection of behavioral health, research, and emerging AI technologies.** He previously directed the Eating Disorder Center of Excellence, where he oversaw multidisciplinary clinical programs, advanced outcomes-driven research, and guided national best-practice development for complex eating-disorder care.
- Dr. Greenberg's career spans federal health systems, large-scale research operations, and digital health strategy. He brings deep clinical training— including pre-doctoral and post-doctoral fellowships in evidence-based treatment, behavioral medicine, and clinical research—paired with executive experience modernizing care models and integrating AI-enabled tools into behavioral health. In his current role with **Clarity Wellness Technologies**, he focuses on responsible AI, clinical decision support, and scalable digital solutions that strengthen access, quality, and patient engagement across the behavioral health continuum.
- Dr. Greenberg completed his PhD in clinical psychology at the Drexel University School of Medicine. He completed his pre-doctoral internship and post-doctoral fellowship at the National Center for PTSD, Boston VAMC with fellowships at Harvard and Boston University Schools of Medicine.

STATEMENT OF INTEREST

- Dr. Greenberg is the Co-Founder and Chief Clinical and Ethics Officer of Clarity Wellness Technologies, Inc. This is an early-stage health technology corporation which applies AI to addressing behavioral health challenges. He holds a financial equity agreement with this organization.

WHAT IS AI: OPINIONS VARY

One End



To the Other

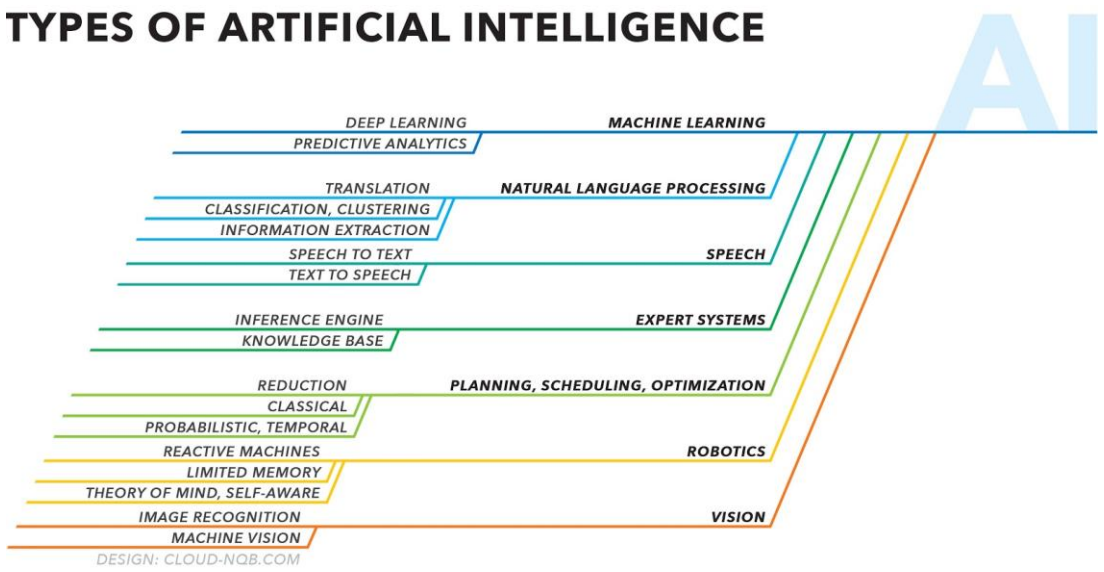


AI: A HISTORY

Alan Turing



Modern AI



AI: A VOCABULARY

- Early AI refers to the first generation of systems built to mimic human reasoning using simple, rule-based logic. These early programs—such as chess computers and mathematical theorem solvers could respond to inputs with predetermined outputs, but they had **no memory, no learning ability, and no awareness of context**. They belong to what is now called **Reactive Machines**, the most basic category of AI, which simply evaluates the current situation and produces a fixed response. Early chess systems like **IBM Deep Blue** exemplify this design, operating entirely through preset rules and pattern evaluation without retaining past experiences .
- As the field progressed, early AI expanded into **Limited Memory systems**, which could use small amounts of past data to improve decision-making. This shift laid the groundwork for modern machine learning. These systems still lacked deeper understanding or self-awareness, but they represented a major step forward by allowing AI to learn from recent experiences—an ability now seen in technologies like self-driving cars . Together, Reactive Machines and Limited Memory AI define the foundational era of artificial intelligence: simple, task-specific systems that established the principles on which today's advanced models are built.

AI: A VOCABULARY

- **Big Data Analytics**
- Big data analytics refers to the advanced methods used to examine extremely large, fast-moving, and diverse datasets to uncover patterns, trends, and actionable insights. These datasets—often too large or complex for traditional data-processing tools require distributed computing, machine learning, and statistical modeling to extract value. In practice, big data analytics enables organizations to make evidence-based decisions, predict future outcomes, personalize services at scale, and optimize operations across domains such as healthcare, finance, and national security.
- At its core, big data analytics integrates four pillars: **volume** (massive amounts of data), **velocity** (rapid generation and processing), **variety** (structured and unstructured data), and **veracity** (data quality and reliability). Techniques such as clustering, predictive modeling, natural language processing, and real-time analytics allow leaders to transform raw information into strategic intelligence. For someone in behavioral health innovation, this means converting millions of clinical, engagement, and environmental signals into insights that improve care quality, access, and outcomes.

AI: A VOCABULARY

- **Neural networks**
- Computational systems inspired by the structure and function of the human brain. At their core, they consist of interconnected layers of nodes—often called **neurons** that process information by passing signals forward through weighted connections. Each neuron receives inputs, multiplies them by learned weights, applies an activation function, and produces an output. When many neurons are layered together, the network can learn complex patterns, relationships, and representations that traditional rule-based systems cannot capture.
- Early neural networks began with the **perceptron**, a single-layer model capable of simple classification. Modern neural networks, however, use multiple layers known as **deep learning** to extract increasingly abstract features from data. This architecture enables breakthroughs in image recognition, natural language processing, speech understanding, and generative AI. In essence, neural networks learn by adjusting their internal weights through exposure to large datasets, allowing them to generalize, predict, and generate new information with remarkable accuracy.

AI: A VOCABULARY

- **Machine learning**
- A branch of artificial intelligence that enables computers to learn patterns from data and improve their performance on a task **without being explicitly programmed**. Instead of following fixed rules, a machine-learning system identifies relationships, makes predictions, or classifies information by adjusting internal parameters based on examples it has seen. This makes it fundamentally different from early rule-based AI, which could only respond with predetermined outputs.
- In practice, machine learning powers many of the AI types shown on your current page—especially **Limited Memory AI**, which learns from past data to make better decisions . Modern ML includes techniques such as supervised learning, unsupervised learning, reinforcement learning, and deep learning, each designed to uncover different kinds of patterns. These methods allow AI systems to recognize images, understand language, forecast trends, and generate new content. Machine learning is the engine behind today's most advanced AI models, including generative systems, predictive analytics, and autonomous agents.

AI: A VOCABULARY

- **Deep learning**
- A specialized branch of machine learning that uses **multi-layered neural networks** to automatically learn complex patterns from large amounts of data. Instead of relying on hand-crafted rules or simple statistical models, deep learning systems extract increasingly abstract features at each layer starting from raw inputs like pixels or text and building up to high-level concepts such as objects, emotions, or semantic meaning. This layered architecture allows deep learning models to outperform traditional approaches in tasks involving vision, speech, natural language, and pattern recognition.
- What makes deep learning powerful is its ability to improve through exposure to massive datasets and adjust millions (or billions) of internal parameters during training. Techniques like convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformers enable breakthroughs in image classification, autonomous driving, language understanding, and generative AI. In short, deep learning is the engine behind today's most advanced AI systems capable of learning representations that were once impossible for machines to capture.

AI: A VOCABULARY

- **Natural language processing (NLP)**
- The field of artificial intelligence focused on enabling computers to understand, interpret, generate, and interact with human language. At its core, NLP transforms raw text or speech into structured representations that machines can work with identifying words, meanings, relationships, and intent. This allows systems to perform tasks such as translation, summarization, sentiment analysis, question answering, and conversational interaction.
- Modern NLP blends linguistics with machine learning, especially neural networks and large language models. These systems learn patterns from massive amounts of text, allowing them to grasp context, nuance, and even the subtleties of tone. NLP is the foundation of chatbots, voice assistants, search engines, and generative AI making human–machine communication more natural, intuitive, and effective.

AI: A VOCABULARY

- **Generative AI**
- The branch of artificial intelligence that **creates new, original content** not just analyzes existing data. It uses machine-learning models, especially deep learning and large language models, to generate text, images, audio, code, and other outputs that did not previously exist. On the page you're viewing, it's defined as AI that *"creates unique new content like text, images, and code"* , and Microsoft further describes it as a subset of AI that uses machine learning to produce new material rather than simply classify or predict.
- At its core, generative AI learns patterns from massive datasets and then uses those patterns to produce novel outputs whether drafting a paragraph, designing a logo, generating synthetic medical data, or writing software. This makes it fundamentally different from predictive AI, which forecasts outcomes based on historical data . Generative AI is now central to modern innovation, powering tools that accelerate creativity, automate complex workflows, and enable more natural human–machine interaction.

AI: A VOCABULARY

- **Predictive AI**
- The category of artificial intelligence that analyzes historical data to **forecast future outcomes, trends, or behaviors**. On the page you're viewing, it's defined as AI that *"forecasts outcomes or trends using historical data"* . These systems don't create new content; instead, they identify patterns in past information and use statistical and machine-learning models to estimate what is likely to happen next.
- Predictive AI powers applications such as risk scoring, demand forecasting, clinical outcome prediction, fraud detection, and personalized recommendations. It relies on techniques like regression, time-series analysis, classification models, and increasingly deep learning. In contrast to generative AI which produces text, images, or code predictive AI is fundamentally about **anticipation**, helping organizations make proactive, data-driven decisions based on modeled probabilities rather than fixed rules.

AI: A VOCABULARY

- **Agentic AI**
- The artificial intelligence systems designed to operate as **autonomous agents** that can *plan, decide, and execute multi-step tasks across different tools, environments, or software systems without needing constant human direction*. On the page you're viewing, it's defined as *"autonomous agents that plan and execute multi-step tasks across systems"* . This distinguishes agentic AI from traditional models that simply generate content or make predictions; agentic systems take action.
- Agentic AI combines reasoning, memory, planning, and tool-use capabilities so an AI can break down a goal into subtasks, choose the right tools or APIs, monitor progress, adapt when conditions change, and complete the objective end-to-end. These systems represent a major evolution in A from passive responders to active problem-solvers enabling workflows like automated research, multi-system coordination, operational optimization, and complex decision support.

AI: A VOCABULARY

- **Artificial general intelligence (AGI)**
- A hypothetical form of AI capable of performing *any* intellectual task a human can—reasoning, learning, adapting, and applying knowledge across domains with true flexibility. Unlike today's systems, which are narrow and task-specific, AGI would demonstrate broad cognitive competence: understanding context, transferring learning between fields, planning long-horizon actions, and operating with human-level autonomy and judgment.
- On the page you're viewing, AGI is described as *"a theoretical form of AI that can perform any intellectual task a human can"* and is noted as **not yet existing** in real-world systems. Current AI—including generative, predictive, and agentic models—remains powerful but fundamentally specialized. AGI represents the next conceptual frontier: machines with generalized intelligence, self-directed learning, and the ability to reason about the world in ways comparable to human cognition.

PERCEIVED BENEFITS OF AI

- 1. Increased Efficiency and Automation
- 2. Enhanced Decision-Making
- 3. Improved Personalization
- 4. Scalability Across Systems
- 5. Better Understanding of Complex Data
- 6. Innovation and Creativity
- 7. Enhanced Human–Machine Interaction
- 8. Continuous Learning and Adaptation

PERCEIVED RISKS OF AI

- 1. Hallucinations and Misinformation
- 2. Bias and Fairness Issues
- 3. Privacy and Data Security Concerns
- 4. Over-reliance on Automation
- 5. Misuse of Generative AI
- 6. Lack of Transparency and Explainability
- 7. Autonomous System Risks
- 8. Societal and Economic Disruption

EARLY SUCCESSES

- 1. Early Detection of Disease (Tumors, Stroke, Cardiac Risk)
- 2. Financial Crime Detection (Money Laundering, Fraud, Illicit Networks)
- 3. Precision Forecasting and Risk Modeling
- 4. Autonomous Systems and Robotics
- 5. Natural Language and Generative AI
- 6. Cybersecurity and Threat Detection

EARLY AI FAILURES

- 1. Medical Misdiagnosis and Unsafe Clinical Recommendations
- 2. Biased or Discriminatory Decision Systems
- 3. Financial AI Failures (Fraud Detection, Trading, AML)
- 4. Safety Failures in Autonomous Vehicles
- 5. Generative AI Misuse and Hallucinations
- 6. AI Agents Taking Unintended Actions