Click to verify

The definition of artificial intelligence goes beyond simple automation – it's the ability of machines to think, learn and adapt. No longer confined to routine tasks, AI now tackles complex challenges once exclusive to human intelligence. It understands language, detects patterns, makes decisions, and even predicts future outcomes with uncanny accuracy. So what can AI do? Today's AI is more powerful than ever. It sees, listens and responds. It learns from experience, refines its skills and integrates seamlessly into our daily lives. From personalized recommendations to fully autonomous systems, AI is transforming the way we innovate, compete and grow in real-time. Self-driving cars? That's just the beginning. AI has crossed a new threshold in the past year. The real game-changer is generative AI – machines that don't just process data, they create. They write code, compose music, generate lifelike images and videos, and even produce entire articles indistinguishable from human work. At the heart of this revolution are machine learning and deep learning, the driving forces accelerating AI's evolution. These technologies are rewriting the rules of innovation, transforming how we interact with technology, and unlocking a future we're only beginning to imagine.What are the benefits of AI? AI technology is redefining how we live and work, driving smarter automation, deeper insights and more strategic decision-making. Here's a look at the key benefits of AI. Automating processes AI takes efficiency to the next level by automating complex workflows and reducing human workload. In cybersecurity, AI-powered systems hunt down threats before they strike. In smart factories, robots with AI-driven vision spot defects, optimize production and keep operations seamless. And companies that use AI in business? They can scale faster, work smarter and do more with less. Zero human error Unlike humans, AI never slips up or gets distracted. It follows strict AI algorithms, ensuring pinpoint accuracy in finance, healthcare and manufacturing. From detecting fraud in banking to perfectly calibrated robotic surgeries, AI enhances reliability across industries. No more repetitive tasks Why waste time on mind-numbing work? AI in business handles document validation, call transcriptions and customer queries – freeing up human talent for creative problem-solving. In hazardous environments, AI-powered robots take over risky jobs, keeping workers safe. Faster, smarter decisions AI processes vast amounts of data at lightning speed, uncovering patterns and insights far beyond human capabilities. It powers real-time financial fraud detection, medical diagnostics and predictive analytics, enabling professionals to stay ahead of the curve. In a world where speed and accuracy are everything, the benefits of AI are game-changing – faster decisions, sharper insights, and the confidence to act before it's too late. 24/7 reliability Forget downtime – AI works around the clock without breaks, fatigue or errors. From cybersecurity monitoring to healthcare diagnostics and customer support, AI technology ensures uninterrupted performance, keeping businesses and services running smoothly around the clock. Accelerating breakthroughs AI is reshaping research and development, driving discoveries in medicine, climate science and engineering. It speeds up drug discovery, deciphers genetic data for personalized medicine and optimizes renewable energy models. With AI, progress happens faster and smarter. How does AI work? At its core, AI processes vast amounts of data, uncovering patterns and making predictions with remarkable precision. It achieves this by leveraging large datasets and intelligent AI algorithms – structured sets of rules that allow software to learn from patterns in the data. The driving force behind this capability? Neural networks: complex systems of interconnected nodes that pass information through multiple layers to find connections and extract meaning from data. To truly understand how AI works, we must unpack the following concepts: Learning: At the heart of AI lies machine learning, enabling systems to analyse data, recognize patterns and make decisions without explicit programming. Taking this further, deep learning uses advanced neural networks to process millions of data points, allowing AI software to understand more complex patterns and continually improve its performance. Reasoning: AI doesn't just recognize trends – it can think and infer. By mimicking human reasoning, AI evaluates commands, context and available data to develop strategies, form hypotheses and make informed decisions in real time. Problem solving: AI approaches problem solving through data manipulation, running simulations, testing different possibilities and refining its strategy. Through intelligent AI algorithms, it explores various possible paths to find the most optimal solution to complex problems. Language processing: AI uses natural language processing – or NLP – to analyse human language data in a way that computers can understand. What is NLP? It refers to the ability of machines to understand, interpret and generate human language, using text analysis, sentiment analysis and machine translation. Perception: Through computer vision, AI-powered systems process data from cameras and sensors to identify objects, detect faces and recognize images with precision. From facial recognition to self-driving cars, perception-driven AI is revolutionizing the way machines interact with the world. Strong AI vs weak AI AI spans a wide spectrum of capabilities, but essentially, it falls into two broad categories: weak AI and strong AI. Weak AI, often referred to as artificial narrow intelligence (ANI) or narrow AI, refers to systems designed to excel at specific tasks within well-defined parameters. These systems operate within a limited scope and lack the capacity for general intelligence. Think of them as highly specialized tools – efficient, precise, but confined to their programmed functions. But don't let the name fool you! Weak AI is anything but weak – it powers countless artificial intelligence applications we interact with daily. Examples of narrow AI are all around us. From Siri and Alexa's instant responses to self-driving cars, ANI is the impetus behind today's most advanced AI innovations. Here are some real-world examples of AI applications powered by narrow AI: Smart assistants: Among the best-known examples of weak AI, digital assistants use natural language processing to handle tasks such as setting reminders, answering questions and controlling smart home devices. Chatbots: Ever chatted with customer support on an e-commerce site? Chances are you were speaking with an ANI-powered chatbot. These AI-driven systems answer routine enquiries, leaving humans free to perform higher-level tasks. Recommendation engines: Whether it's Netflix curating your next must-watch series or Amazon predicting your next purchase, ANI analyses user habits to provide personalized recommendations based on viewing, buying or browsing patterns. Navigation apps: How do you get from A to B without getting lost? Apps like Google Maps rely on ANI algorithms to process real-time traffic data, optimize routes and guide users to their destinations efficiently. Email spam filters: Do you wonder why most spam emails never reach your inbox? ANI-powered filters scan messages, detect suspicious content and redirect unwanted emails to the spam folder. Autocorrect features: Whether you're texting on an iPhone or composing an email, AI software refines your writing by correcting typos and suggesting words based on your typing patterns, ensuring smoother, more efficient communication. Each of these applications showcases ANI's ability to tackle specific tasks by leveraging large datasets and specialized algorithms. So, the next time you're impressed by AI's capabilities, remember – it's weak AI driving these remarkable innovations, transforming our world in ways we once thought impossible. Also known as artificial general intelligence (AGI) Designed to adapt, learn and apply knowledge across various domains Also known as artificial narrow intelligence (ANI) or narrow AI Designed to excel at specific tasks within well-defined parameters In contrast to narrow AI, the concept of strong AI – also known as general AI – aims to develop systems capable of handling a broad range of tasks with human-like proficiency. Unlike their task-specific ANI counterparts, strong AI systems aspire to possess a form of general intelligence that enables them to learn, adapt and apply knowledge across multiple domains. The ultimate goal? To create artificial entities with cognitive abilities that mirror those of humans, capable of engaging in intellectual tasks spanning diverse fields. For now, strong AI remains purely speculative, with no practical examples in real life. However, that hasn't stopped AI researchers from pushing the boundaries of AI's potential development. Research in artificial general intelligence (AGI) is exploring how AI could evolve beyond its specialized functions into autonomous systems capable of independent reasoning. In theory, AGI could take on any human job, whether it's cleaning, coding or scientific research. While we're not there yet, the potential impact of AGI spans multiple industries, including: Language: Writing essays, poems and engaging in conversations. Healthcare: Medical imaging, drug research and surgery. Transportation: Fully automated cars, trains and planes. Arts and entertainment: Creating music, visual art and films. Domestic robots: Cooking, cleaning and childcare. Manufacturing: Supply chain management, stocktaking and consumer services. Engineering: Programming, building and architecture. Security: Detecting fraud, preventing security breaches and improving public safety. While researchers and developers continue to push the limits of AGI, achieving true general intelligence, on a par with human cognition, remains a formidable challenge and a distant goal. That said, with rapid advancements in AI technology and machine learning, the real question is no longer if AGI will emerge, but when.What are the four types of AI? Artificial intelligence spans a wide range of capabilities, each designed for specific functions and objectives. Understanding the four types of AI provides insight into the ever-evolving landscape of machine intelligence. Reactive machines: These AI systems operate strictly within predefined rules but lack the ability to learn from new data or experiences. A prime example is chatbots, which generate responses based on programmed algorithms, rather than adapting to conversations. While they excel at specific tasks, they cannot evolve beyond their initial programming. Limited memory: Unlike reactive machines, AI systems with limited memory can learn from historical data, enabling them to make informed decisions based on past experiences. These types of AI are seen in self-driving cars, which use sensors and machine learning algorithms to analyse traffic patterns and navigate safely through dynamic environments. Similarly, natural language processing applications leverage historical data to refine language comprehension and interpretation over time. Theory of mind: Still theoretical, this type of AI would be capable of understanding human emotions, intentions and social cues. A machine with a theory of mind could then use that information to anticipate human actions and engage in intuitive, empathetic interactions. If realized, this AI could revolutionize human-computer and social robotics, creating systems that genuinely understand us. Self-aware AI: The most futuristic (and controversial) concept, self-aware AI refers to machines with human-like consciousness – aware of their own existence and capable of perceiving emotions in others. While captivating in science-fiction classics like Blade Runner, this level of AI remains purely hypothetical, sparking both fascination and debate about the future of artificial intelligence. These four types of AI highlight the vast spectrum of intelligence within artificial systems. As AI technology advances, exploring the capabilities and limitations of each type will deepen our understanding of machine intelligence and its impact on society.Machine learning vs deep learning Central to these advancements are machine learning and deep learning, two subfields of AI that drive many of today's innovations. While closely related, each has its own distinct approach to learning and problem solving. Machine learning relies on different learning methods to train AI systems. The three primary types are: Supervised learning: The algorithm is trained on a labelled dataset, where each input has a known output. By learning from these labelled examples, the model can make accurate predictions on new, unseen data. Unsupervised learning: Unlike supervised learning, this method works without predefined labels or outputs. Instead, the algorithm learns to identify hidden structures or groupings within the data, making it essential for tasks like clustering or anomaly detection. Reinforcement learning: In this approach, an AI agent interacts with an environment and learns through trial and error. It receives rewards for desirable actions or penalties for mistakes, gradually improving its decision-making over time. This technique is widely used in robotics, gaming and autonomous systems. A subset of machine learning, deep learning focuses on training artificial neural networks with multiple layers, inspired by the human brain's structure and function. These networks consist of interconnected nodes (neurons) that process and transmit signals, enabling AI to learn complex patterns. Unlike traditional machine learning models, deep learning algorithms automatically extract features from raw data, refining their understanding through layers of abstraction. This makes them exceptionally powerful in image and speech recognition, natural language processing and other advanced AI applications. Yet their high complexity comes at a cost – deep learning requires massive datasets, extensive training and significant computational power to achieve optimal performance.Examples of AI technology While many people associate AI with smart assistants like Siri and Alexa, new AI technology is emerging fast, making daily tasks more efficient and transforming industries in unexpected ways. Here are some key applications: Healthcare: AI can process and analyse vast amounts of patient data, enabling accurate diagnoses, predictive analytics and personalized treatment recommendations for better health outcomes. It also plays a crucial role in drug discovery and medical imaging, helping doctors detect diseases earlier and more effectively. Business and manufacturing: AI-driven automation enhances efficiency across industries, from fraud detection and risk assessment to market trend analysis. In manufacturing, AI-powered robots streamline production while predictive maintenance helps prevent equipment failures before they happen. In retail, AI enables personalized shopping experiences, smart inventory management, chatbots for customer support and data-driven advertising strategies to increase sales. Education: AI-powered intelligent tutoring systems adapt to students' learning styles, providing personalized feedback and guidance. AI also automates grading, content creation and virtual-reality simulations, making education more interactive and efficient. Transport: AI keeps traffic moving, prevents breakdowns and streamlines logistics in shipping and supply chains. From fleet tracking to automated scheduling, it ensures faster, smarter and more efficient operations. Agriculture: AI-driven drones and sensors monitor soil health, detect crop diseases and optimize irrigation. Smart systems also recommend efficient pesticide use and resource management, helping farmers maximize crop yields with minimal waste. Entertainment: AI curates personalized recommendations, matching you with the perfect movie, song or book based on your preferences. Virtual and augmented reality push immersion to new levels, while AI-driven CGI and special effects bring movies and games to life with stunning realism. The growth and impact of generative AI The rise of large-scale language models like Chat GPT is just the beginning. Welcome to the era of generative AI – a groundbreaking frontier in artificial intelligence that goes beyond analysing data to creating entirely new content. Unlike traditional AI systems, which excel at classification and prediction, generative models push boundaries by mimicking human creativity and imagination. They generate text, images, music, and even entire virtual worlds, blurring the line between machine output and human innovation. But generative AI isn't flawless. While its capabilities are revolutionary, challenges remain. Deepfakes, misinformation, biases, copyright issues and job displacement are all real concerns. These generative models also demand immense computational power, driving up costs and environmental impact while posing security and quality control risks. Despite these hurdles, examples of artificial intelligence in this space continue to expand, proving its extraordinary potential. Researchers are actively tackling these challenges through improved detection systems, refined training data, enhanced security measures and optimized computational efficiency. A balanced approach, supported by guidelines and stronger regulation, will also be key to ensuring generative AI serves as a force for progress, not disruption.AI governance and regulations As AI becomes deeply embedded in industries worldwide, ensuring the quality and reliability of AI software is more critical than ever. Yet, despite its rapid growth, AI still operates in a largely unregulated space, posing risks that demand urgent attention. This is where International Standards come in. Standards, such as those developed by ISO/IEC JTC 1/SC 42 on artificial intelligence, play a pivotal role in addressing the responsible development and use of AI technologies. They provide decision makers and policymakers with a structured framework to create consistent, auditable and transparent AI systems, closing regulatory gaps. For businesses, aligning with these standards isn't just about compliance – it's a strategic advantage. From risk management to responsible AI governance, standardized AI practices enhance credibility, build trust with stakeholders, and ensure that the benefits of artificial intelligence outweigh the risks.History of artificial intelligence: who invented AI? AI has progressed in leaps and bounds, transforming many aspects of our world. But to truly appreciate its current capabilities, it's important to understand its origins and evolution. So who created AI? To find out, let's take a journey through the fascinating history of artificial intelligence. Today's AI loosely stems from the 19th-century invention of Charles Babbage's "difference engine" – the world's first successful automatic calculator. British code-breaker Alan Turing, who was a key figure in the Allies' intelligence arsenal during WWII, amongst other feats, can also be seen as a father figure of today's iterations of AI. In 1950, he proposed the Turing Test, designed to assess a machine's ability to exhibit intelligent behaviour indistinguishable from that of a human. From that point onward, advancements in AI technology began to accelerate exponentially, spearheaded by such influential figures as John McCarthy, Marvin Minsky, Herbert Simon, Geoffrey Hinton, Yoshua Bengio, Yann LeCun, and many others. But it wasn't all smooth sailing. While AI flourished in the early years, with computers' capability to store more information, it soon hit a roadblock: computers simply couldn't store enough information or process it fast enough. It wasn't until the 1980s that AI experienced a renaissance, sparked by an expansion of the algorithm toolkit and an increase in funding. To cut a long story short, here are some key events and milestones in the history of artificial intelligence: 1950: Alan Turing publishes the paper "Computing Machinery and Intelligence", in which he proposes the Turing Test as a way of assessing whether or not a computer counts as intelligent. 1956: A small group of scientists gather for the Dartmouth Summer Research Project on Artificial Intelligence, which is regarded as the birth of this field of research. 1966-1974: This is conventionally known as the "First AI Winter", a period marked by reduced funding and progress in AI research due to failure to live up to early hype and expectations. 1997: Deep Blue, an IBM chess computer, defeats world champion Garry Kasparov in a highly publicized chess match, demonstrating the fabulous potential of AI systems. In the same year, speech recognition software, developed by Dragon Systems, was implemented on Windows. 2011: In a televised Jeopardy! contest, IBM's Watson Deep QA computer defeats two of the quiz shows' all-time champions, showcasing the ability of AI systems to understand natural language. 2012: The "deep learning" approach, inspired by the human brain, revolutionizes many AI applications, ushering in the current AI boom. 2016: Developed by a Google subsidiary, the computer program AlphaGo captures the world's attention when it defeats legendary Go player Lee Sedol. The ancient board game "Go" is one of the most complex ever created. 2017 to date: Rapid advancements in computer vision, natural language processing, robotics and autonomous systems are driven by progress in deep learning and increased computational power. 2023: The rise of large language models, such as GPT-3 and its successors, demonstrates the potential of AI systems to generate human-like text, answer questions and assist with a wide range of tasks. 2024: New breakthroughs in multimodal AI allow systems to process and integrate various types of data (text, images, audio and video) for more comprehensive and intelligent solutions. AI-powered digital assistants are now capable of engaging in natural, contextual conversations as well as assisting with a wide variety of tasks. How will AI change our world? The exponential growth of computing power and the Internet has propelled machine learning from concept to reality. Today, AI algorithms don't just follow instructions, they learn from vast datasets, improving with each iteration. At its most advanced, this has led to deep learning, where computers refine their "intelligence" through experience, much like the human brain. And the impact? AI is everywhere – powering how we work, communicate and engage with technology. From medical breakthroughs to climate solutions, its impact will be profound and far-reaching. But with innovation comes responsibility. As AI becomes more powerful and pervasive, we must ensure it is developed and used responsibly. For this to be achieved, it is crucial to stay informed and be proactive in shaping its development – to build a future that is both beneficial and empowering for all. Use your Google Account to sign in to DeepAI Artificial intelligence (AI) refers to computer systems capable of performing complex tasks that historically only a human could do, such as reasoning, making decisions, or solving problems. Today, the term "AI" describes a wide range of technologies that power many of the services and goods we use every day – from apps that recommend TV shows to chatbots that provide customer support in real time. But do all of these really constitute artificial intelligence as most of us envision it? And if not, then why do we use the term so often? In this article, you'll learn more about artificial intelligence, what it actually does, and different types of it. In the end, you'll also learn about some of its benefits and dangers and explore flexible courses that can help you expand your knowledge of AI even further. Enroll in AI for Everyone, an online program offered by DeepLearning.AI. In just 6 hours, you'll gain foundational knowledge about AI terminology, strategy, and the workflow of machine learning projects. Your first week is free. What is artificial intelligence?Artificial intelligence (AI) is the theory and development of computer systems capable of performing tasks that historically required human intelligence, such as recognizing speech, making decisions, and identifying patterns. AI is an umbrella term that encompasses a wide variety of technologies, including machine learning, deep learning, and natural language processing (NLP). Although the term is commonly used to describe a range of different technologies in use today, many disagree on whether these actually constitute artificial intelligence. Instead, some argue that much of the technology used in the real world today actually constitutes highly advanced machine learning that is simply a first step towards true artificial intelligence, or "general artificial intelligence" (GAI).Yet, despite the many philosophical disagreements over whether "true" intelligent machines actually exist, when most people use the term AI today, they're referring to a suite of machine learning-powered technologies, such as Chat GPT or computer vision, that enable machines to perform tasks that previously only humans can do like generating written content, steering a car, or analyzing data. Read more: The History of AI: A Timeline of Artificial IntelligenceAI stands for "artificial intelligence." Artificial intelligence is the simulation of human intelligence processes by machines, such as computer systems. AI powers many technology-driven industries, such as health care, finance, transportation, and much more.Artificial intelligence examples Though the humanoid robots often associated with AI (think Star Trek: The Next Generation's Data or Terminator's T-800) don't exist yet, you've likely interacted with machine learning-powered services or devices many times before. At the simplest level, machine learning uses algorithms trained on data sets to create machine learning models that allow computer systems to perform tasks like making song recommendations, identifying the fastest way to travel to a destination, or translating text from one language to another. Some of the most common examples of AI in use today include: ChatGPT: Uses large language models (LLMs) to generate text in response to questions or comments posed to it. Google Translate: Uses deep learning algorithms to translate text from one language to another. Netflix: Uses machine learning algorithms to create personalized recommendation engines for users based on their previous viewing history. Tesla: Uses computer vision to power self-driving features on their cars. Read more: Deep Learning vs. Machine Learning: Beginner's GuideThe increasing accessibility of generative AI tools has made it an in-demand skill for many tech roles. If you're interested in learning to work with AI for your career, you might consider a free, beginner-friendly online program like Google's Introduction to Generative AI. AI in the workforceArtificial intelligence is prevalent across many industries. Automating tasks that don't require human intervention saves money and time, and can reduce the risk of human error. Here are a couple of ways AI could be employed in different industries:Finance industry. Fraud detection is a notable use case for AI in the finance industry. AI's capability to analyze large amounts of data enables it to detect anomalies or patterns that signal fraudulent behavior. Health care industry. AI-powered robotics could support surgeries close to highly delicate organs or tissue to mitigate blood loss or risk of infection. Subscribe to our weekly newsletter Career Chat. It's a low-commitment way to stay current with industry trends and skills you can use to guide your career path. What is artificial general intelligence (AGI)? Artificial general intelligence (AGI) refers to a theoretical state in which computer systems will be able to achieve or exceed human intelligence. In other words, AGI is "true" artificial intelligence, as depicted in countless science fiction novels, television shows, movies, and comics. As for the precise meaning of "AI" itself, researchers don't quite agree on how we would recognize "true" artificial general intelligence when it appears. However, the most famous approach to identifying whether a machine is intelligent or not is known as the Turing Test or Imitation Game, an experiment that was first outlined by influential mathematician, computer scientist, and cryptanalyst Alan Turing in a 1950 paper on computer intelligence. There, Turing described a three-player game in which a human "interrogator" is asked to communicate via text with another human and a machine and judge who composed each response. If the interrogator cannot reliably identify the human, then Turing says the machine can be said to be intelligent [1]. To complicate matters, researchers and philosophers also can't quite agree whether we're beginning to achieve AGI, if it's still far off, or just totally impossible. For example, while a recent paper from Microsoft Research and OpenAI argues that Chat GPT-4 is an early form of AGI, many other researchers are skeptical of these claims and argue that they were just made for publicity [2, 3].Regardless of how far we are from achieving AGI, you can assume that when someone uses the term artificial general intelligence, they're referring to the kind of sentient computer programs and machines that are commonly found in popular science fiction. Read more: Artificial General Intelligence vs. AIWhen researching artificial intelligence, you might have come across the terms "strong" and "weak" AI. Though these terms might seem confusing, you likely already have a sense of what they mean. Strong AI is essentially AI that is capable of human-level, general intelligence. In other words, it's just another way to say "artificial general intelligence." Weak AI, meanwhile, refers to the narrow use of widely available AI technology, like machine learning or deep learning, to perform very specific tasks, such as playing chess, recommending songs, or steering cars. Also known as Artificial Narrow Intelligence (ANI), weak AI is essentially the kind of AI we use daily.Read more: Machine Learning vs. AI: Differences, Uses, and BenefitsThe 4 Types of AI As researchers attempt to build more advanced forms of artificial intelligence, they must also begin to formulate more nuanced understandings of what intelligence or even consciousness precisely mean. In their attempt to clarify these concepts, researchers have outlined four types of artificial intelligence.Here's a summary of each AI type, according to Professor Arend Hintze of the University of Michigan [4]: 1. Reactive machinesReactive machines are the most basic type of artificial intelligence. Machines built in this way don't possess any knowledge of previous events but instead only "react" to what is before them in a given moment. As a result, they can only perform certain advanced tasks within a very narrow scope, such as playing chess, and are incapable of performing tasks outside of their limited context. 2. Limited memory machinesMachines with limited memory possess a limited understanding of past events. They can interact more with the world around them than reactive machines can. For example, self-driving cars use a form of limited memory to make turns, observe approaching vehicles, and adjust their speed. However, machines with only limited memory cannot form a complete understanding of the world because their recall of past events is limited and only used in a narrow band of time. 3. Theory of mind machinesMachines that possess a "theory of mind" represent an early form of artificial general intelligence. In addition to being able to create representations of the world, machines of this type would also have an understanding of other entities that exist within the world. As of this moment, this reality has still not materialized. 4. Self-aware machinesMachines with self-awareness are the theoretically most advanced type of AI and would possess an understanding of the world, others, and itself. This is what most people mean when they talk about achieving AGI. Currently, this is a far-off reality. Generative AI is a kind of artificial intelligence capable of producing original content, such as written text or images, in response to user inputs or "prompts." Generative models are also known as large language models (LLMs) because they're essentially complex, deep learning models trained on vast amounts of data that can be interacted with using normal human language rather than technical jargon. Generative AI is becoming increasingly common in everyday life, powering tools such as ChatGPT, Google Gemini, and Microsoft Copilot. While other kinds of machine learning models are well suited for performing narrow, repetitive tasks, generative AI is capable of responding to user inputs with unique outputs that allow it to respond dynamically in real-time. This makes it particularly useful for powering interactive programs like virtual assistants, chatbots, and recommendation systems. That said, while generative AI may produce responses that make it seem like self-aware AI, the reality is that its responses are the result of statistical analysis rather than sentience. Read more: Generative AI Examples and How the Technology WorksAI benefits and dangersAI has a range of applications with the potential to transform how we work and live. While many of these transformations are exciting, like self-driving cars, virtual assistants, or wearable devices in the healthcare industry, they also pose many challenges.It's a complicated picture that often summons competing images: a utopia for some, a dystopia for others. The reality is likely to be much more complex. Here are a few of the possible benefits and dangers AI may pose: Potential BenefitsPotential DangersGreater accuracy for certain repeatable tasks, such as assembling vehicles or computers.Job loss due to increased automation.Decreased operational costs due to greater efficiency of machines.Potential for bias or discrimination as a result of the data set on which the AI is trained.Increased personalization within digital services and products.Possible cybersecurity concerns.Improved decision-making in certain situations.Lack of transparency over how decisions are arrived at, resulting in less than optimal solutions.Ability to quickly generate new content, such as text or images.Potential to create misinformation, as well as inadvertently violate laws and regulations.These are just some of the ways that AI provides benefits and dangers to society. When using new technologies like AI, it's best to keep a clear mind about what it is and isn't. With great power comes great responsibility, after all. Read more: AI Ethics: What It Is and Why It MattersBuild AI skills on CourseraArtificial Intelligence is quickly changing the world we live in. If you're interested in learning more about AI and how you can use it at work or in your own life, consider taking one of these courses or specializations on Coursera today: For a quick overview of AI, take DeepLearning.AI's AI For Everyone course. There, you'll learn what AI can realistically do and not do, how to spot opportunities to apply AI to problems in your own organization, and what it feels like to build machine learning and data science projects. Top build job-ready AI skills to enhance your career, enroll in the IBM AI Foundations for Everyone Specialization. Learn foundational AI concepts, explore AI tools and services, and engage with AI environments through hands-on projects. To learn how AI can address complex real-world problems, explore DeepLearning.AI's AI For Good Specialization. Here, you'll build skills combining human and machine intelligence for positive real-world impact using AI in a beginner-friendly, three-course program. Technology Computers Ask the Chatbot a Question artificial intelligence (AI), the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings. The term is frequently applied to the project of developing systems endowed with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience. Since their development in the 1940s, digital computers have been programmed to carry out very complex tasks—such as discovering proofs for mathematical theorems or playing chess—with great proficiency. Despite continuing advances in computer processing speed and memory capacity, there are as yet no programs that can match full human flexibility over wider domains or in tasks requiring much everyday knowledge. On the other hand, some programs have attained the performance levels of human experts and professionals in executing certain specific tasks, so that artificial intelligence in this limited sense is found in applications as diverse as medical diagnosis, computer search engines, voice or handwriting recognition, and chatbots. All but the simplest human behavior is ascribed to intelligence, while even the most complicated insect behavior is usually not taken as an indication of intelligence. What is the difference? Consider the behavior of the digger wasp, Sphex ichneumoneus. When the female wasp returns to her burrow with food, she first deposits it on the threshold, checks for intruders inside her burrow, and only then, if the coast is clear, carries her food inside. The real nature of the wasp's instinctual behavior is revealed if the food is moved a few inches away from the entrance to her burrow while she is inside: on emerging, she will repeat the whole procedure as often as the food is displaced. Intelligence—conspicuously absent in the case of the wasp—must include the ability to adapt to new circumstances. Psychologists generally characterize human intelligence not by just one trait but by the combination of many diverse abilities. Research in AI has focused chiefly on the following components of intelligence: learning, reasoning, problem solving, perception, and using language. There are a number of different forms of learning as applied to artificial intelligence. The simplest is learning by trial and error. For example, a simple computer program for solving mate-in-one chess problems might try moves at random until mate is found. The program might then store the solution with the position so that, the next time the computer encountered the same position, it would recall the solution. This simple memorizing of individual items and procedures—known as rote learning—is relatively easy to implement on a computer. More challenging is the problem of implementing what is called generalization. Generalization involves applying past experience to analogous new situations. For example, a program that learns the past tense of regular English verbs by rote will not be able to produce the past tense of a word such as jump unless the program was previously presented with jumped, whereas a program that is able to generalize can learn the "add -ed" rule for regular verbs ending in a consonant and so form the past tense of jump on the basis of experience with similar verbs. (Read Ray Kurzweil's Britannica essay on the future of "Nonbiological Man.") Computers and Technology Quiz Artificial intelligence (AI) is technology that enables computers and machines to simulate human learning, comprehension, problem solving, decision making, creativity and autonomy. Applications and devices equipped with AI can see and identify objects. They can understand and respond to human language. They can learn from new information and experience. They can make detailed recommendations to users and experts. They can act independently, replacing the need for human intelligence or intervention (a classic example being a self-driving car). But in 2024, most AI researchers, practitioners and most AI-related headlines are focused on breakthroughs in generative AI (gen AI), a technology that can create original text, images, video and other content. To fully understand generative AI, it's important to first understand the technologies on which generative AI tools are built: machine learning (ML) and deep learning. A simple way to think about AI is as a series of nested or derivative concepts that have emerged over more than 70 years: Directly underneath AI, we have machine learning, which involves creating models by training an algorithm to make predictions or decisions based on data. It encompasses a broad range of techniques that enable computers to learn from and make inferences based on data without being explicitly programmed for specific tasks. There are many types of machine learning techniques or algorithms, including linear regression, logistic regression, decision trees, random forest, support vector machines (SVMs), k-nearest neighbor (KNN), clustering and more. Each of these approaches is suited to different kinds of problems and data. But one of the most popular types of machine learning algorithm is called a neural network (or artificial neural network). Neural networks are modeled after the human brain's structure and function. A neural network consists of interconnected layers of nodes (analogous to neurons) that work together to process and analyze complex data. Neural networks are well suited to tasks that involve identifying complex patterns and relationships in large amounts of data. The simplest form of machine learning is called supervised learning, which involves the use of labeled data sets to train algorithms to classify data or predict outcomes accurately. In supervised learning, humans pair each training example with an output label. The goal is for the model to learn the mapping between inputs and outputs in the training data, so it can predict the labels of new, unseen data. Deep learning is a subset of machine learning that uses multilayered neural networks, called deep neural networks, that more closely simulate the complex decision-making power of the human brain. Deep neural networks include an input layer, at least three but usually hundreds of hidden layers, and an output layer, unlike neural networks used in classic machine learning models, which usually have only one or two hidden layers. These multiple layers enable unsupervised learning: they can automate the extraction of features from large, unlabeled and unstructured data sets, and make their own predictions about what the data represents. Because deep learning doesn't require human intervention, it enables machine learning at a tremendous scale. It is well suited to natural language processing (NLP), computer vision, and other tasks that involve the fast, accurate identification complex patterns and relationships in large amounts of data. Some form of deep learning powers most of the artificial intelligence (AI) applications in our lives today. Deep learning also enables: Semi-supervised learning, which combines supervised and unsupervised learning by using both labeled and unlabeled data to train AI models for classification and regression tasks. Self-supervised learning, which generates implicit labels from unstructured data, rather than relying on labeled data sets for supervisory signals. Reinforcement learning, which learns by trial-and-error and reward functions rather than by extracting information from hidden patterns. Transfer learning, in which knowledge gained through one task or data set is used to improve model performance on another related task or different data set. Generative AI, sometimes called "gen AI", refers to deep learning models that can create complex original content such as long-form text, high-quality images, realistic video or audio and more in response to a user's prompt or request. At a high level, generative models encode a simplified representation of their training data, and then draw from that representation to create new work that's similar, but not identical, to the original data. Generative models have been used for years in statistics to analyze numerical data. But over the last decade, they evolved to analyze and generate more complex data types. This evolution coincided with the emergence of three sophisticated deep learning model types: Variational autoencoders or VAEs, which were introduced in 2013, and enabled models that could generate multiple variations of content in response to a prompt or instruction. Diffusion models, first seen in 2014, which add "noise" to images until they are unrecognizable, and then remove the noise to generate original images in response to prompts. Transformers (also called transformer models), which are trained on sequenced data to generate extended sequences of content (such as words in sentences, shapes in an image, frames of a video or commands in software code). Transformers are at the core of most of today's headline-making generative AI tools, including ChatGPT and GPT-4, Copilot, BERT, Bard and Midjourney. In general, generative AI operates in three phases: Training, to create a foundation model.Tuning, to adapt the model to a specific application.Generation, evaluation and more tuning, to improve accuracy. Generative AI begins with a "foundation model"; a deep learning model that serves as the basis for multiple different types of generative AI applications. The most common foundation models today are large language models (LLMs), created for text generation applications. But there are also foundation models for image, video, sound or music generation, and multimodal foundation models that support several kinds of content. To create a foundation model, practitioners train a deep learning algorithm on huge volumes of relevant raw, unstructured, unlabeled data, such as terabytes or petabytes of data text or images or video from the internet. The training yields a neural network of billions of parameters encoded representations of the entities, patterns and relationships in the data that can generate content autonomously in response to prompts. This is the foundation model. This training process is compute-intensive, time-consuming and expensive. It requires thousands of clustered graphics processing units (GPUs) and weeks of processing, all of which typically costs millions of dollars. Open source foundation model projects, such as Meta's Llama-2, enable gen AI developers to avoid this step and its costs. Next, the model must be tuned to a specific content generation task. This can be done in various ways, including: Fine-tuning, which involves feeding the model application-specific labeled data, questions or prompts the application is likely to receive, and corresponding correct answers in the wanted format. Reinforcement learning with human feedback (RLHF), in which human users evaluate the accuracy or relevance of model outputs so that the model can improve itself. This can be as simple as having people type or talk back corrections to a chatbot or virtual assistant. Developers and users regularly assess the outputs of their generative AI apps, and further tune the model even as often as once a week for greater accuracy or relevance. In contrast, the foundation model itself is updated much less frequently, perhaps every year or 18 months. Another option for improving a gen AI app's performance is retrieval augmented generation (RAG), a technique for extending the foundation model to use relevant sources outside of the training data to refine the parameters for greater accuracy or relevance. An AI agent is an autonomous AI program, it can perform tasks and accomplish goals on behalf of a user or another system without human intervention, by designing its own workflow and using available tools (other applications or services). Agentic AI is a system of multiple AI agents, the efforts of which are coordinated, or orchestrated, to accomplish a more complex task or a greater goal than any single agent in the system could accomplish. Unlike chatbots and other AI models which operate within predefined constraints and require human intervention, AI agents and agentic AI exhibit autonomy, goal-driven behavior and adaptability to changing circumstances. The terms "agent" and "agentic" refer to these models' agency, or their capacity to act independently and purposefully. One way to think of agents is as a natural next step after generative AI. Gen AI models focus on creating content based on learned patterns; agents use that content to interact with each other and other tools to make decisions, solve problems and complete tasks. For example, a gen AI app might be able to tell you the best time to climb Mt. Everest given your work schedule, but an agent can tell you this, and then use an online travel service to book you the best flight and reserve a room in the most convenient hotel in Nepal. AI offers numerous benefits across various industries and applications. Some of the most commonly cited benefits include: Automation of repetitive tasks.More and faster insight from data. Enhanced decision-making. Fewer human errors.24x7 availability.Reduced physical risks. AI can automate routine, repetitive and often tedious tasks including digital tasks such as data collection, entering and preprocessing, and physical tasks such as warehouse stock-picking and manufacturing processes. This automation frees to work on higher value, more creative work. Whether used for decision support or for fully automated decision-making, AI enables faster, more accurate predictions and reliable, data-driven decisions. Combined with automation, AI enables businesses to act on opportunities and respond to crises as they emerge, in real time and without human intervention. AI can reduce human errors in various ways, from guiding people through the proper steps of a process, to flagging potential errors before they occur, and fully automating processes without human intervention. This is especially important in industries such as healthcare where, for example, AI-guided surgical robotics enable consistent precision. Machine learning algorithms can continually improve their accuracy and further reduce errors as they're exposed to more data and "learn" from experience. AI is always on, available around the clock, and delivers consistent performance every time. Tools such as AI chatbots or virtual assistants can lighten staffing demands for customer service or support. In other applications such as materials processing or production lines, AI can help maintain consistent work quality and output levels when used to complete repetitive or tedious tasks. By automating dangerous work such as animal control, handling explosives, performing tasks in deep ocean water, high altitudes or in outer space, AI can eliminate the need to put human workers at risk of injury or worse. While they have yet to be perfected, self-driving cars and other vehicles offer the potential to reduce the risk of injury to passengers. The real-world applications of AI are many. Here is just a small sampling of use cases across various industries to illustrate its potential: Companies can implement AI-powered chatbots and virtual assistants to handle customer inquiries, support tickets and more. These tools use natural language processing (NLP) and generative AI capabilities to understand and respond to customer questions about order status, product details and return policies. Chatbots and virtual assistants enable always-on support, provide faster answers to frequently asked questions (FAQs), free human agents to focus on higher-level tasks, and give customers faster, more consistent service. Machine learning and deep learning algorithms can analyze transaction patterns and flag anomalies, such as unusual spending or login locations, that indicate fraudulent transactions. This enables organizations to respond more quickly to potential fraud and limit its impact, giving themselves and customers greater peace of mind. Retailers, banks and other customer-facing companies can use AI to create personalized customer experiences and marketing campaigns that delight customers, improve sales and prevent churn. Based on data from customer purchase history and behaviors, deep learning algorithms can recommend products and services customers are likely to want, and even generate personalized copy and special offers for individual customers in real time. AI-driven recruitment platforms can streamline hiring by screening resumes, matching candidates with job descriptions, and even conducting preliminary interviews using video analysis. These and other tools can dramatically reduce the mountain of administrative paperwork associated with fielding a large volume of candidates. It can also reduce response times and time-to-hire, improving the experience for candidates whether they get the job or not. Generative AI code generation tools and automation tools can streamline repetitive coding tasks associated with application development, and accelerate the migration and modernization (reformatting and replatorming) of legacy applications at scale. These tools can speed up tasks, help ensure code consistency and reduce errors. Machine learning models can analyze data from sensors, Internet of Things (IoT) devices and operational technology (OT) to forecast when maintenance will be required and predict equipment failures before they occur. AI-powered preventive maintenance helps prevent downtime and enables you to stay ahead of supply chain issues before they affect the bottom line. Organizations are scrambling to take advantage of the latest AI technologies and capitalize on AI's many benefits. This rapid adoption is necessary, but adopting and maintaining AI workflows comes with challenges and risks. AI systems rely on data sets that might be vulnerable to data poisoning, data tampering, data bias or cyberattacks that can lead to data breaches. Organizations can mitigate these risks by protecting data integrity and implementing security and availability throughout the entire AI lifecycle, from development to training and deployment and postdeployment. Threat actors can target AI models for theft, reverse engineering or unauthorized manipulation. Attackers might compromise a model's integrity by tampering with its architecture, weights or parameters; the core components that determine a model's behavior, accuracy and performance. Like all technologies, models are susceptible to operational risks such as model drift, bias and breakdowns in the governance structure. Left unaddressed, these risks can lead to system failures and cybersecurity vulnerabilities that threat actors can use. If organizations don't prioritize safety and ethics when developing and deploying AI systems, they risk committing privacy violations and producing biased outcomes. For example, biased training data used for hiring decisions might reinforce gender or racial stereotypes and create AI models that favor certain demographic groups over others. AI ethics is a multidisciplinary field that studies how to optimize AI's beneficial impact while reducing risks and adverse outcomes. Principles of AI ethics are applied through a system of AI governance consisted of guardrails that help ensure that AI tools and systems remain safe and ethical. AI governance encompasses oversight mechanisms that address risks. An ethical approach to AI governance requires the involvement of a wide range of stakeholders, including developers, users, policymakers and ethicists, helping to ensure that AI-related systems are developed and used to align with society's values. Here are common values associated with AI ethics and responsible AI: Explainability and interpretability As AI becomes more advanced, humans are challenged to comprehend and retrace how the algorithm came to a result. Explainable AI is a set of processes and methods that enables human users to interpret, comprehend and trust the results and output created by algorithms. Fairness and inclusion Although machine learning, by its very nature, is a form of statistical discrimination, the discrimination becomes objectionable when it places privileged groups at systematic advantage and certain unprivileged groups at systematic disadvantage, potentially causing varied harms. To encourage fairness, practitioners can try to minimize algorithmic bias across data collection and model design, and to build more diverse and inclusive teams. Robustness and security Robust AI effectively handles exceptional conditions, such as abnormalities in input or malicious attacks, without causing unintentional harm. It is also built to withstand intentional and unintentional interference by protecting against exposed vulnerabilities. Accountability and transparency Organizations should implement clear responsibilities and governance structures for the development, deployment and outcomes of AI systems. In addition, users should be able to see how an AI service works, evaluate its functionality, and comprehend its strengths and limitations. Increased transparency provides information for AI consumers to better understand how the AI model or service was created. Privacy and compliance Many regulatory frameworks, including GDPR, mandate that organizations abide by certain privacy principles when processing personal information. It is crucial to be able to protect AI models that might contain personal information, control what data goes into the model in the first place, and to build adaptable systems that can adjust to changes in regulation and attitudes around AI ethics. In order to contextualize the use of AI at various levels of complexity and sophistication, researchers have defined several types of AI that refer to its level of sophistication: Weak AI: Also known as "narrow AI," defines AI systems designed to perform a specific task or a set of tasks. Examples might include "smart" voice assistant apps, such as Amazon's Alexa, Apple's Siri, a social media chatbot or the autonomous vehicles promised by Tesla. Strong AI: Also known as "artificial general intelligence" (AGI) or "general AI," possess the ability to understand, learn and apply knowledge across a wide range of tasks at a level equal to or surpassing human intelligence. This level of AI is currently theoretical and no known AI systems approach this level of sophistication. Researchers argue that if AGI is even possible, it requires major increases in computing power. Despite recent advances in AI development, self-aware AI systems of science fiction remain firmly in that realm. The idea of "a machine that thinks" dates back to ancient Greece. But since the advent of electronic computing (and relative to some of the topics discussed in this article) important events and milestones in the evolution of AI include the following: 1950 Alan Turing publishes Computing Machinery and Intelligence. In this paper, Turing famous for breaking the German ENIGMA code during WWII and often referred to as the "father of computer science" asks the following question: "Can machines think?" From there, he offers a test, now famously known as the "Turing Test," where a human interrogator would try to distinguish between a computer and human text response. While this test has undergone much scrutiny since it was published, it remains an important part of the history of AI, and an ongoing concept within philosophy as it uses ideas around linguistics. 1956 John McCarthy coins the term "artificial intelligence" at the first-ever AI conference at Dartmouth College. (McCarthy went on to invent the Lisp language.) Later that year, Allen Newell, J.C. Shaw and Herbert Simon create the Logic Theorist, the first-ever running AI computer program. 1967 Frank Rosenblatt builds the Mark 1 Perceptron, the first computer based on a neural network that "learned" through trial and error. Just a year later, Marvin Minsky and Seymour Papert publish a book titled Perceptrons, which becomes both the landmark work on neural networks and, at least for a while, an argument against future neural network research initiatives. 1980 Neural networks, which use a backpropagation algorithm to train itself, became widely used in AI applications. 1995 Stuart Russell and Peter Norvig publish Artificial Intelligence: A Modern Approach, which becomes one of the leading textbooks in the study of AI. In it, they delve into four potential goals or definitions of AI, which differentiates computer systems based on rationality and thinking versus acting. 1997 IBM's Deep Blue beats then world chess champion Garry Kasparov, in a chess match (and rematch). 2004 John McCarthy writes a paper, What Is Artificial Intelligence?, and proposes an often-cited definition of AI. By this time, the era of big data and cloud computing is underway, enabling organizations to manage ever-larger data estates, which will one day be used to train AI models. 2011 IBM Watson® beats champions Ken Jennings and Brad Rutter at Jeopardy! Also, around this time, data science begins to emerge as a popular discipline. 2015 Baidu's Minwa supercomputer uses a special deep neural network called a convolutional neural network to identify and categorize images with a higher rate of accuracy than the average human. 2016 DeepMind's AlphaGo program, powered by a deep neural network, beats Lee Sodol, the world champion Go player, in a five-game match. The victory is significant given the huge number of possible moves as the game progresses (over 14.5 trillion after just four moves). Later, Google purchased DeepMind for a reported USD 400 million. 2022 A rise in large language models or LLMs, such as OpenAI's ChatGPT, creates an enormous change in performance of AI and its potential to drive enterprise value. With these new generative AI practices, deep-learning models can be pretrained on large amounts of data. 2024 The latest AI trends point to a continuing AI renaissance. Multimodal models that can take multiple types of data as input are providing richer, more robust experiences. These models bring together computer vision image recognition and NLP speech recognition capabilities. Smaller models are also making strides in an age of diminishing returns with massive models with large parameter counts. Related solutions IBM watsonx.ai Train, validate, tune and deploy generative AI, foundation models and machine learning capabilities with IBM watsonx.ai, a next-generation enterprise studio for AI builders. Build AI applications in a fraction of the time with a fraction of the data. Discover watsonx.ai Artificial intelligence solutions Put AI to work in your business with IBM's industry-leading AI expertise and portfolio of solutions at your side. Explore AI solutions AI consulting and services Reinvent critical workflows and operations by adding AI to maximize experiences, real-time decision-making and business value. Explore AI services Get one-stop access to capabilities that span the AI development lifecycle. Produce powerful AI solutions with user-friendly interfaces, workflows and access to industry-standard APIs and SDKs. Explore watsonx.ai Book a live demo