



CHEAT SHEET:

ARTIFICIAL INTELLIGENCE

by Aminu Abdullahi

ARTIFICIAL INTELLIGENCE

AMINU ABDULLAHI

Artificial intelligence comes in many forms, from simple tools that respond to customers via a chat to complex [machine learning](#) algorithms that predict the trajectory of an entire organization. Despite years of overpromising, AI doesn't comprise sentient machines that reason like humans. Rather, AI encompasses more narrowly focused pattern matching at scale to complement human reasoning.

In order to help business leaders understand what AI capabilities are, how to use artificial intelligence and where to begin an AI journey, it's essential to first dispel the myths surrounding this huge leap in AI technology.

What is artificial intelligence?

AI is largely a pattern-recognition tool that can run at a scale that's dramatically beyond any human, yet never quite replaces humans. Even at its best, AI delivers acceptable, though not perfect results, giving people the ability to step in, observe the data and reason from there.

Note that while we use AI throughout this cheat sheet, most enterprises actually engage with a subset of AI called machine learning or [deep learning](#). We'll use AI here as a shorthand that includes machine learning and deep learning.

The truth is that current AI technology is limited, but it's still incredibly powerful. However complicated its processes may seem in practice, at the core of AI-driven applications is the simple ability to [identify patterns and make inferences](#) based on those patterns.

AI isn't truly intelligent, and it's often as biased as the data we choose to feed into our ML models. That doesn't mean AI isn't useful for businesses and consumers trying to solve real-world problems, it means that we're nowhere close to machines that can actually make independent decisions or arrive at conclusions without being given the proper data first. It's also true that AI can tend to confirm our biases, rather than eliminate them.

How does artificial intelligence work?

AI is a complex system designed to model human behavior and intelligence. It combines large data with intelligent algorithms to analyze, understand, and make decisions or predictions about future states. To make accurate predictions, AI systems require large amounts of data to learn from; this data is gathered from various sources, processed, analyzed and organized in a suitable format for the AI algorithms.

AI algorithms are the core of AI systems and are designed to analyze and interpret data, identify patterns, and make predictions or decisions based on the input. By continuously collecting new data and retraining the models, AI systems can adapt to changing conditions and improve their performance.

The core process of how AI works involves the following subdomains:

- **Machine learning:** A branch of AI that focuses on the development of algorithms and statistical models that allow computer systems to learn and improve from data without being explicitly programmed.
- **Deep learning:** A subfield of machine learning that mimics the workings of the human brain's neural networks, using multiple layers of artificial neural networks to learn and understand complex patterns and features in data.
- **Neural networks:** A computational model, inspired by the structure and function of the human brain, that can process and analyze large amounts of data to recognize patterns, make predictions or classify information.
- **Natural language processing:** A branch of AI that focuses on the interaction between computers and human language, enabling machines to understand, interpret and generate human language.
- **Computer vision:** A branch of AI that enables machines to interpret and understand visual information from images or videos.
- **Cognitive computing:** A model that aims to create AI systems that can simulate human-like intelligence and interact with humans in a more natural and intuitive way.

What can artificial intelligence do?

Artificial intelligence is essentially pattern matching at scale. With its pattern recognition capabilities, modern AI can perform image recognition, understand the natural language and writing patterns of humans, make connections between different types of data, identify abnormalities in patterns, strategize, predict and more.

While humans are unable to as easily comb through the amount of data that machines can to uncover patterns, machines struggle when presented with an outlier that might be easy for a human to spot but contradicts the training data. Therefore, the best AI applications are highly focused and combine human reasoning with the brute power of ML.

Since the onset of the COVID-19 pandemic in 2020, AI and ML have seen a [massive market growth](#). The global pandemic also shifted AI priorities and applications: Instead of solely focusing on financial analysis and consumer insight, post-pandemic [AI projects have trended toward customer experience](#) and cost optimization.

AI bots can perform many basic customer service tasks, freeing employees up to only address cases that need human intervention. AI like this has been particularly widespread since the start of the pandemic, when workers forced out of call centers put stress on customer service.

What are the business applications of AI?

In the business world, there are plenty of AI applications, but perhaps none is gaining traction as much as business and predictive analytics and its end goal: [prescriptive analytics](#).

Business analytics is a complicated set of processes that aim to model the present state of a business, predict where it will go if kept on its current trajectory and model potential futures with a given set of changes. Predicting the future with an established model of the past can be easy enough, but prescriptive analysis, which aims to find the best possible outcome by tweaking an organization's current course, can be downright impossible without AI help.

Analytics may be the rising star of business AI, but it's hardly the only application of artificial intelligence in the commercial and industrial worlds. Other AI use cases for businesses include the following:

- **Recruiting and employment:** AI can streamline recruiting by filtering through larger numbers of candidates more quickly than a human and by noticing qualified people who may be overlooked.

- **Fraud detection:** AI is great at picking up on subtle differences and irregular behavior, such as subtle indicators of financial fraud that humans may miss.
- **Cybersecurity:** AI is great at detecting indicators of hacking and other [cybersecurity issues](#).
- **Data management:** Using AI, you can categorize raw data and find relationships between items that were previously unknown.
- **Customer relations:** Modern AI-powered chatbots are incredibly good at carrying on conversations thanks to natural language processing, making them a great first line of customer service.
- **Healthcare:** Not only are some AI applications able to detect cancer and other health concerns before doctors, they can also provide feedback on patient care based on long-term records and trends.
- **Predicting market trends:** Much like prescriptive analysis in the business analytics world, AI systems can be trained to predict trends in larger markets, which can lead to businesses getting a jump on emerging trends.
- **Reducing energy use:** AI can streamline energy use in buildings and even across cities as well as make better predictions for construction planning, oil and gas drilling, and other [energy-centric projects](#).
- **Marketing:** AI systems can be trained to increase the value of marketing both toward individuals and larger markets, helping organizations save money and get better marketing results.

What are the different types of AI?

Narrow AI

Also known as weak AI, narrow AI helps you perform specific functions. It's focused on a single domain and operates within predefined limits. Narrow AI cannot do anything more than what they are programmed to do — they have a very limited or narrow range of competencies. Examples include voice assistants like Siri or [Alexa](#). These lack general intelligence and cannot perform tasks outside their designated domain.

General AI

General AI, also known as strong AI or artificial general intelligence, refers to AI systems that possess human-level intelligence and can understand, learn and perform any intellectual task that a human being can do.

They can adapt to various scenarios and solve problems creatively. While general AI remains a long-term objective, current advancements primarily focus on developing narrow AI systems that excel in specific areas, such as image recognition or natural language processing.

Superintelligent AI

This type of AI surpasses human intelligence in nearly all aspects. It not only outperforms humans in cognitive tasks but also possesses the ability to improve itself, leading to an exponential increase in intelligence. While superintelligent AI remains largely theoretical at present, it's a topic of interest and concern within the field of AI.

Reactive machines

Reactive AI systems automatically respond to a limited set or combination of inputs and operate based on the current input without any memory or past experiences. In fact, reactive AI systems don't have the ability to form memories or learn from previous interactions. They simply react to the current situation or stimulus. Examples include chess-playing AI systems or recommendation algorithms.

Limited memory

AI systems with limited memory can store and retrieve information from previous experiences to make better decisions. They have the ability to learn from past data and use it to improve their future actions. Self-driving cars that use historical data to make driving decisions are an example of limited memory AI.

Theory of mind

This type of AI is still largely theoretical. Theory of mind AI will have the ability to understand and model the mental states, beliefs, and intentions of other agents. They'll be able to attribute thoughts, emotions, and intentions to other entities and predict their behavior based on this understanding.

Self-aware

While still theoretical and not fully realized, this type of AI would possess human-like awareness and understanding of its own existence. Self-aware AI systems will have a sense of their own existence, consciousness and internal state. They possess self-reflective abilities and are aware of their own thoughts, actions and impact on their environment. True self-aware AI is a concept that remains largely speculative and is currently beyond the capabilities of existing technology.

Generative AI

Generative AI systems are capable of creating content, such as images, videos, music or text, that is nearly indistinguishable from human-generated content. They can autonomously generate new outputs based on their learned patterns and styles.

Generative adversarial networks are an example of generative AI, where one network generates content, and another network evaluates and provides feedback to improve the quality of the generated output.

Other popular examples of generative AI include:

- **DeepArt.io:** This tool uses neural networks to transform photos into artistic styles from famous artists.
- **Runway:** This platform offers a range of generative AI tools for creating images and videos.
- **DeepDream:** DeepDream is a tool developed by Google that uses generative AI to modify images. It enhances patterns and structures in an image to create dream-like visuals.
- **OpenAI's ChatGPT:** Generative pretrained transformer is a language model developed by OpenAI to generate human-like text based on a given prompt. The latest model GPT-4 was trained on Microsoft Azure AI supercomputers and is available on ChatGPT Plus.

What AI platforms are available?

When adopting an AI strategy, it's important to know what software is available for business-focused AI. There are a wide variety of platforms available from the usual cloud computing suspects like Google, AWS, Microsoft and IBM, and choosing the right one can mean the difference between success and failure.

AWS Machine Learning offers a wide variety of services that run in the AWS cloud. AI services, prebuilt frameworks, analytics tools and more are all available, with many designed to take the legwork out of getting started and others like SageMaker for Business Analysts designed to enable corporations to get AI insight without writing code. AWS offers prebuilt AI algorithms, one-click ML training and training tools for developers getting started in or expanding their knowledge of AI development.

Google Cloud offers similar AI solutions to AWS, as well as having several prebuilt total AI solutions that organizations can ideally plug into their organizations with minimal effort. Google also distinguishes itself by innovating some of the industry standards for AI like TensorFlow, an open-source ML library.

Microsoft's AI platform comes with pre-generated services, ready-to-deploy cloud computing infrastructure and a variety of additional AI tools that can be plugged into existing models. Its AI Lab also offers a wide range of AI apps that developers can tinker with and learn from what others have done. Microsoft also offers an AI school with educational tracks specifically for business applications.

Watson is IBM's version of cloud-hosted ML and business AI, but it goes a bit farther with more AI options. IBM offers on-site servers custom built for AI tasks for businesses that don't want to rely on cloud hosting, and it also has IBM AI OpenScale, an AI platform that can be integrated into other cloud hosting services, which could help to avoid vendor lock-in. In 2021, IBM Watson suffered a **media backlash** after years of overpromising on what its AI could deliver in healthcare, but many enterprises still turn to it for narrower tasks.

What AI skills will businesses need to invest in?

Perhaps the most critical skill needed to use AI is knowing when to skip AI altogether. The reality of AI is that many problems could be solved by applying simple regression analysis or if/then statements. Most AI, in other words, isn't AI at all: It's only basic math and common sense.

For more complicated, AI-oriented tasks, the associated data science breaks down into two categories: that which is intended for human consumption and that which is intended for machine consumption.

In the latter case, AI involves complex digital models that apply ML models and AI algorithms to large amounts of data. These systems then act autonomously to generate a particular ad or customer experience or make real-time stock trades. Hence, machine-oriented AI professions require "exceptionally strong mathematical, statistical and computational fluency to build models that can quickly make good predictions," as former Google and Foursquare data scientist Michael Li has [noted](#).

By contrast, the skills needed for more human-oriented data science and AI skew toward storytelling. Given that no data is unbiased, the role of human intelligence is to help the data tell clear stories. Such AI storytellers use data visualization to facilitate exploration and insights into that data.

For many in AI, the most sophisticated math they'll do is power analyses and significance tests. They might write SQL queries to get data, do basic math on that data, graph results and then explain the results. Not gee whiz data science, but it's incredibly helpful for breaking down complex data into actionable insights, to use the data science lingo.

With all that in mind, it's still true that skills needed for an AI project differ based on business needs and the platform being used, though most of the biggest platforms support most, if not all, of the most commonly used AI programming languages and skills needed.

Many business AI platforms offer training courses in the specifics of running their architecture and the programming languages needed to develop more AI tools. Businesses that are serious about AI should plan to either hire new employees or give existing ones the [time and resources](#) necessary to train in the skills needed to make AI projects succeed.

How can businesses start using AI?

Getting started with business AI isn't as easy as simply spending money on an AI platform provider and spinning up some prebuilt models and algorithms. There's a lot that goes into successfully adding AI to an organization.

At the heart of it all is [good project planning](#). Adding artificial intelligence to a business, no matter how it will be used, is like any business transformation initiative. Here's an outline of just one way to approach getting started with business AI.

Determine your AI objective

To begin, figure out how AI can be used in your organization and to what end. By focusing on a narrower implementation with a specific goal, you can better allocate resources.

Identify what needs to happen to get there

Once you know where you want to be, you can figure out where you are and how to make the journey. This could include starting to sort existing data, gathering new data, hiring talent and other preproject steps.

Build a team

With an end goal in sight and a plan to get there, it's time to assemble the best team to make it happen. This can include current employees, but don't be afraid to go outside the organization to find the most qualified people. Be sure to allow existing staff to train so they have the opportunity to contribute to the project.

Choose an AI platform

Some AI platforms may be better suited to particular projects, but by and large they all offer similar products in order to compete with each other. Let your team give recommendations on which AI platform to choose — they're the experts who will be in the trenches.

Begin implementation

With a goal, team and platform, you're ready to start working in earnest. This won't be quick: AI machines need to be trained, testing on subsets of data has to be performed and lots of tweaks will need to be made before a business AI is ready to hit the real world. In fact, you should expect that the **vast majority of your time** won't be spent in crafting sexy algorithms, but rather in data preparation.