eazystock.com

What is artificial intelligence?

An overview of artificial intelligence, its uses and benefits for inventory management

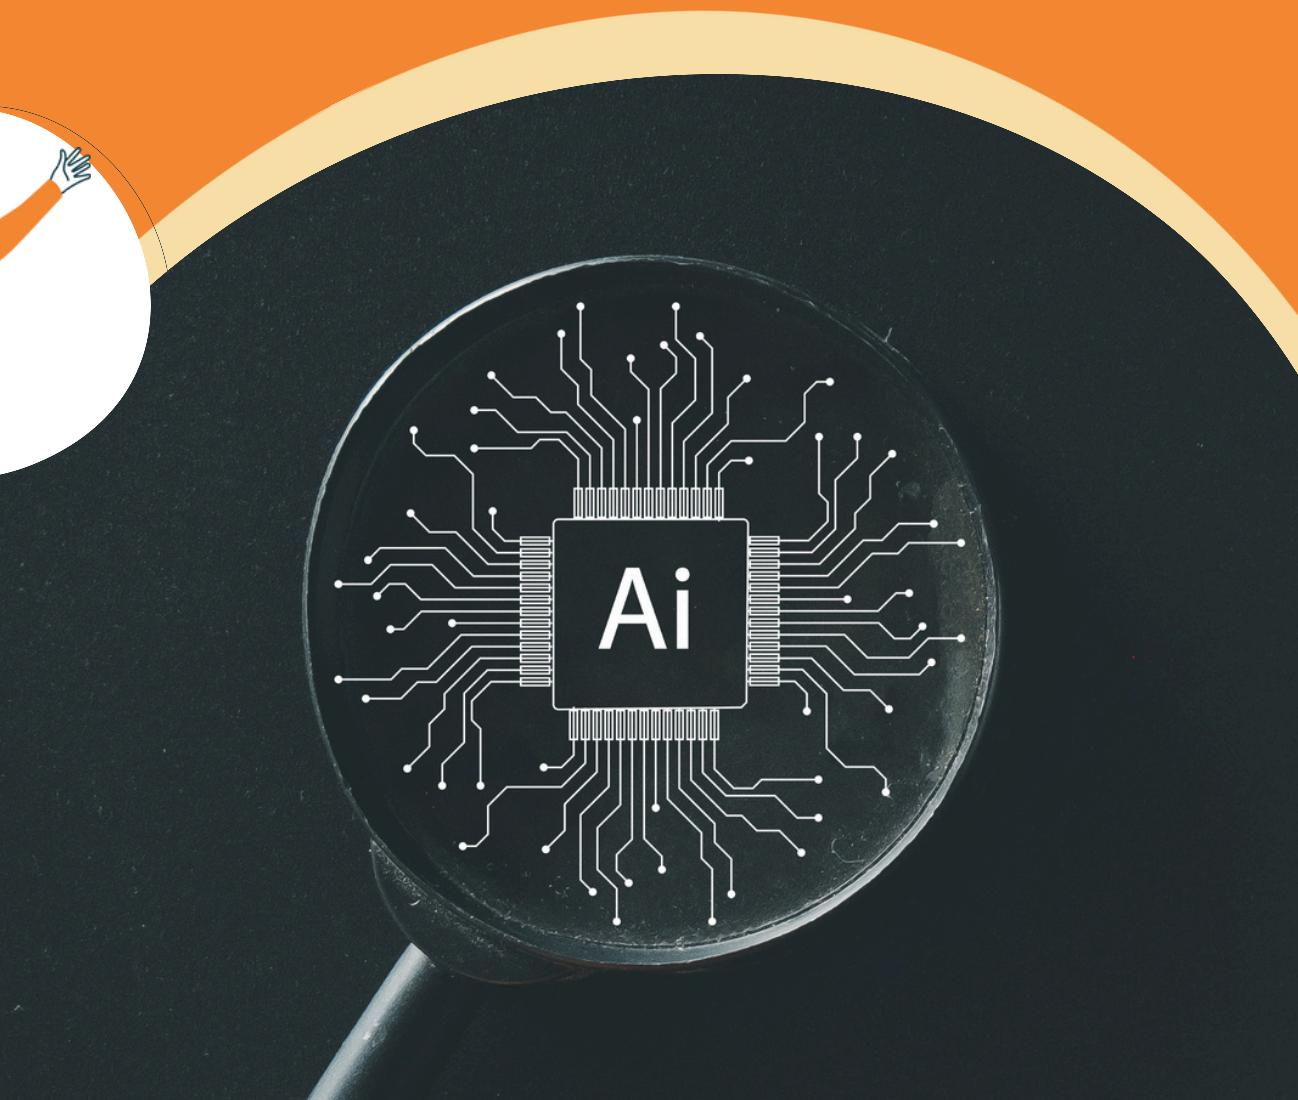

Contents

An explanation of artificial intelligence

AI technology is developing rapidly, with new features hitting the headlines almost daily as companies jostle to showcase their products in an increasingly competitive market.

You can't move for adverts promoting AI-enabled devices, systems, and software, telling us how much they will make our lives easier. Still, there's more to artificial intelligence (AI) than editing photos on our phones and drafting an email. AI is helping to solve global problems, such as reducing emissions and improving accessibility, more efficiently than humans can.

That's not all. AI is being used to analyse data across healthcare to identify patterns that could support new medicines, cures, and preventative methods. It can also highlight areas for improvement to streamline transportation and save time and money, such as minimising wheel friction and finding efficient driving routes.

AI can improve manufacturing efficiency by using robots and analysing production lines to identify bottlenecks and better working methods.

Even as we write this article, Word suggests words and sentences using artificial intelligence to make writing easier.

While there are many positives to AI, there are also some negatives. In the wrong hands, AI can mislead and manipulate people's opinions. For example, by using AI-generated content that uses fake imagery, videos and voices to create false information and deepfakes, which are then shared online and through social media.

This BBC article explains more about deepfake technology, how deepfakes work, and how to spot them.

With social manipulation, if someone views false or harmful information, AI algorithms learn what they have seen and will keep sharing similar fake information. The more false information a user sees, the more they will believe and act on it, as was the case in the UK with recent riots.

In an effort to address concerns and risks of implementing AI, the European Commission enacted the European Artificial Intelligence Act (AI Act) on 1 August 2024. The AI Act introduces a framework for EU countries to follow using a risk-based approach.

With so many areas to consider, there's a lot to take in and understand about AI. In this article, we will break down some AI basics to make them easier to understand so you can decide how to use it to improve your work and home lives.

What is artificial intelligence?

There are many similar definitions of artificial intelligence, so we asked ChatGPT, a generative AI[1] chatbot and virtual assistant to explain. (OpenAI. (2024). ChatGPT (4o) [Large language model]. https://chatgpt.com)

Artificial intelligence (AI) refers to the simulation of human intelligence in machines that are programmed to think and learn like humans. These systems are designed to perform tasks that typically require human intelligence, such as visual perception, speech recognition, decision-making, and language translation.

AI comprises six subsets: machine learning, deep learning, natural language processing, expert systems, robotics, machine vision, and speech recognition. (CourseRA)

[1] Generative AI learns from large online datasets like text and images to generate new content that looks and feels like it was created by a human. It can also generate videos and music that mimics other artists.

How does artificial intelligence work?

Our expertise lies in inventory optimisation, so we're going to hand it over to How Stuff Works to explain how artificial intelligence works.

"AI systems work by combining large amounts of data with intelligent algorithms, which are a series of instructions that allow the software to learn from patterns and features of the data."

SAS expands on this by saying, *"AI works by combining large amounts of data with fast, iterative processing and intelligent algorithms, allowing the software to learn automatically from patterns or features in the data."*

How does artificial intelligence learn?

AI systems learn by processing enormous amounts of data to identify patterns to model decision-making. According to aimunch.com, there are three main types of AI learning – supervised, unsupervised and reinforcement.

We've summarised these below, along with three others – semi-supervised, transfer and self-supervised:

Supervised learning

With supervised learning, the AI system is trained on labelled data sets and parameters. Labelled datasets provide context so the model can answer correctly.

Supervised learning models focus on learning input and output data relationships. For example, AI could be trained to classify images of different stock items like screws, basins or other products.

Human interaction in supervised learning can reinforce good decisions and discourage bad ones to ensure false information doesn't get through.

Unsupervised learning

Unsupervised learning is where algorithms learn raw data's inherent structure without specific guidance or instruction. As the data doesn't have labels, the algorithm identifies naturally occurring patterns in the dataset.

An example of unsupervised learning is analysing sales data to highlight anomalies so they can be investigated and demand forecasts adjusted accordingly.

Reinforcement learning

Reinforcement learning is similar to human learning with positive and negative reinforcement. AI learns to make and improve decision-making by interacting with its environment and receiving rewards or penalties based on those decisions. This is used in robotics and autonomous vehicles. For example, if you use autonomous vehicles in a warehouse, they will be rewarded for collecting the correct items and reaching their destination. However, they will be penalised for collisions or not reaching the end warehouse destination.

As www.that-ai.com explains, reinforcement learning is different to supervised and unsupervised as the AI system creates its own data by being an active part of the environment of rewards and punishments.

Semi-supervised learning

With semi-supervised learning, AI models use a combination of supervised and unsupervised learning, so they are trained on labelled and unlabelled data. Some examples of semi-supervised learning are speech recognition, web content classification and text document classification, as Altexsoft explains.

Transfer learning

As the name suggests, transfer learning reduces the need to build a model from scratch by sharing logic or algorithms between systems. Some real-world examples, as Packt shows, include transferring knowledge from a robot trained using simulation to another robot.

Self-supervised learning

As with unsupervised learning, self-supervised learning uses raw data without manual labelling. The model trains itself to create labels, learn different input parts and get the correct answers from the data.

Self-supervised learning is used in natural language processing, machine vision, speech recognition, healthcare and robotics. Most generative AI models like ChatGPT use self-supervised learning. You can read more about self-supervised learning on Grammarly.

What can artificial intelligence be used for?

As mentioned earlier, AI has six subsets with a wide range of uses. The main ones are machine learning and deep learning. The others – natural language processing, expert systems, robotics, machine vision and speech recognition – are machine learning and deep learning subsets.

Here, we explain the different types of AI and provide examples of how they are used.

Machine learning

Machine learning is when machines use experiences to learn without any explicit programming. Instead, they design and develop algorithms and statistical models so they can perform complex tasks. Machine learning uses supervised learning, unsupervised learning and reinforcement learning.

Computer systems use machine learning to process large amounts of historical data to identify patterns. As the system accesses more data, it develops and enhances the algorithm to improve performance and accuracy.

There are multiple steps in the machine-learning process, as laid out by Data Camp.

1. Data collection – collecting data relevant to your problem in a suitable format.
2. Data preprocessing – cleaning data, handling missing data and ensuring it's in a standard format.
3. Choosing the right model – decide based on the size and type of your data, the problem you're trying to solve and the computational resources available.
4. Training the model – feed the data into the model and allow it to adjust its internal parameters.
5. Evaluating the model – test the model on new data it didn't see during training. Check for accuracy, precision and recall, and mean squared error.
6. Hyperparameter tuning and optimisation – you could try different parameter combinations and cross-validation to finetune.
7. Predictions and deployment - feed new data into the model and use the model's output to help decision-making.

There are two fundamental components of machine learning: labels and features.

Labels are the desired outcomes or predictions we want to make, while features are the measurable characteristics or attributes of the data that help us make those predictions.

The relationship between labels and features allows for extracting meaningful patterns from the data. Models can then be trained more effectively, leading to more accurate predictions on new, unseen data. Careful consideration and thoughtful engineering of features ensure that the model captures the relevant information and generalises well.

Data Economy Solutions expand on labels and features in their article.

Inventory example

For example, EazyStock's AI-powered inventory optimisation software uses advanced artificial intelligence, machine learning and optimisation technology to improve demand forecasts accuracy to minimise capital investment in unnecessary inventory, and boost product availability.

EazyStock's machine learning algorithms consider historical sales and purchasing data, trends, seasonality, and product lifecycles to avoid stockouts, minimise waste, and cut unnecessary supply chain costs.

The system can analyse more data than is humanly possible and quickly identify patterns or problems in the data and build relationships between changes in features such as prices, web-search volumes and product groups so you can adapt your inventory plans to reflect demand changes before they cause problems.

Machine learning can help with predictive maintenance. Predictive maintenance uses historical data, such as previous service dates and failures, to anticipate future problems. This saves repair costs and ensures your warehouses and production lines endure minimum downtime.

With smart buildings and warehouses, AI can monitor real-time energy consumption and weather conditions to make data-driven, energy-optimisation decisions. This reduces energy wastage and saves money.

The Internet of Things (IoT) allows businesses to connect items using sensors and tracking devices like NFC, QR codes, barcodes and tags to track inventory items, their condition and expected delivery times. This enables quick response to issues, reduces the risk of loss and saves tracking administration time.

Deep learning

Deep learning is part of machine learning, so it also uses structured, unstructured, and reinforcement learning. The main difference is that deep learning can learn from its mistakes, while machine learning needs human intervention.

For example, if a machine learning model makes a mistake, human intervention is required to adjust it. With a deep-learning model, an algorithm can determine whether a prediction is accurate through its own neural network – no human help is needed.

In deep learning, algorithms are structured in layers to create an artificial neural network that mimics the human brain so a machine can carry out tasks without human intervention.

Deep learning algorithms are complex, as IBM explains. They explain six different types of neural networks to address specific problems or datasets:

1. Convolutional neural networks (CNNs or ConvNets) – primarily used in computer vision and image classification applications.
2. Recurrent neural networks (RNNs) – used in natural language and speech recognition applications.
3. Autoencoders and variational autoencoders – the first deep-learning models widely used for generating realistic images and speech.
4. Generative adversarial networks (GANs) – used in and outside AI to create new data resembling the original training data.
5. Diffusion models – generate data similar to the data they were trained on, then overwrite it.
6. Transformer models – using fill-in-the-blank guessing, an encoder learns the relationship between words and sentences to build up a powerful representation of language without labelling parts of speech and other grammatical features.

Deep learning combines other machine learning techniques to teach algorithms to build models that can be used in generative AI applications like ChatGPT. Self-driving cars and speech and image recognition also use deep learning.

Natural language processing

Natural language processing (NLP) is when a computer program can understand and communicate spoken and written human language.

Thanks to a combination of computational linguistics, statistical modelling, and machine and deep learning, computers can recognise, understand, and generate text and speech.

Large language models LLMs are deep learning algorithms that use supervised, unsupervised and reinforcement learning to perform natural language processing tasks.

Inventory example

EazyStock’s AI help assistant lets users ask questions about their data and the system. It recognises and understands the question and searches through learned data to provide accurate answers. The user can then delve further into the subject with additional questions or using questions suggested by the help assistant.

NLP can also support your customer service teams by using chatbots to answer initial customer enquiries based on knowledge. The chatbot will escalate conversations to a customer service team member when they cannot answer questions directly. This frees up team members’ time to deal with the most critical questions to help keep customers happy.

Expert systems

Expert systems use machine learning to emulate human experts' decision-making skills from their knowledge. These systems don't use conventional procedural code. Instead, they use knowledge sets, rules and reasoning to solve complex problems that usually require a human expert.

The three main uses of expert systems are medical diagnosis, financial planning and troubleshooting. Expert systems like NHS Direct can help diagnose diseases and recommend treatments based on patient symptoms and medical history. For financial planning, expert systems can provide personalised financial advice and investment strategies based on individual goals and risk tolerance.

Robotics

Robots have been used for decades without AI. They have been programmed to carry out repetitive tasks and move heavy objects without thinking

Combining robotics with AI and engineering significantly increases possible applications for robots to carry out tasks automatically or semi-automatically, predominantly through reinforcement learning. Robots can perform repetitive and dangerous tasks and those requiring heavy lifting faster, more efficiently, and safely than humans. The most common use is in factories on production lines.

Inventory example

Robots can be beneficial for picking items in large warehouses, especially if there is a labour shortage. Using robots for repetitive tasks can increase efficiency and accuracy and reduce the risk of injuries.

Another form of robotics is automated vehicles, which are more suitable for smaller warehouses. They can improve warehouse efficiency by planning the shortest and most efficient routes to move items around warehouses, particularly where items are moving along regular routes.

Analysing robotics and order data means inventory items can be placed in the most appropriate place depending on their pick frequency and value to the business.

Machine vision

Machine vision is a subset of machine learning and deep learning.

It uses optic sensors enabled with machine learning tools so machines can capture and analyse visual information to recognise objects using structured learning techniques. Once the machine captures an image, a processor analyses it using AI and machine learning algorithms. It then sends the information back to the machine to make an appropriate decision. Cameras and digital signal processing allow the machines to count items and read serial numbers.

Inventory example

Machine-vision-enabled machines like drones can be used for stock counting, picking orders and warehouse management. [PD1] For example, drones can fly around a warehouse and scan items using sensors, then send the data to the ERP or warehouse management system. This is especially useful for items stacked high. Drones can be programmed with the warehouse layout so they can design efficient paths and avoid crashes. They can also be used to ensure pallets and boxes are stored correctly in the warehouse.

Machine vision can also be used for quality control in manufacturing by examining items in more detail than humans to pick up faults faster and more efficiently.

Speech recognition

Speech recognition is another subset of machine learning that uses natural language processing (NLP).

It enables machines to understand speech and commands to carry out tasks. This common AI feature is used daily on smartphones and smart speakers to ask questions or play music. It's also used in cars to set navigation systems or make hands-free phone calls.

Inventory example

Instead of using handheld devices and scanning items, warehouse staff can wear natural-language-processing-enabled headsets and microphones connected to a mobile computer. The voice system can advise teams on which items need picking and packing. Users can also confirm they've completed a task or ask questions directly.

Will artificial intelligence replace humans?

There's been much controversy about using AI to replace human roles. While AI can perform some roles more efficiently than humans, it won't replace them. However, those who don't adopt AI-enabled technology are more likely to be replaced by those who do.

Forbes believes there will be significant shifts in the job market as AI advances. They also say that while AI might replace certain roles, it brings new opportunities in sectors needing skills that AI cannot replicate, such as complex decision-making, emotional intelligence and creative skills. Understanding these trends allows companies to plan education and training to help people transition to roles where human expertise remains irreplaceable.

Professor Michael Bronstein from the University of Oxford is optimistic about the capabilities of AI technologies to transform society and improve our lives. He imagines AI will augment human capabilities rather than replace them.

While AI might replace some roles, it brings opportunities to enhance roles. For example, AI-enabled technology can improve efficiency and free up time to focus on more strategic tasks.

ChaptGPT has become renowned for drafting content and is embedded in social media tools and platforms to offer AI post creation. However, because these systems learn by processing and reading other websites, they will need checking and editing to avoid plagiarism, ensure they make sense, and have the proper context for your needs.

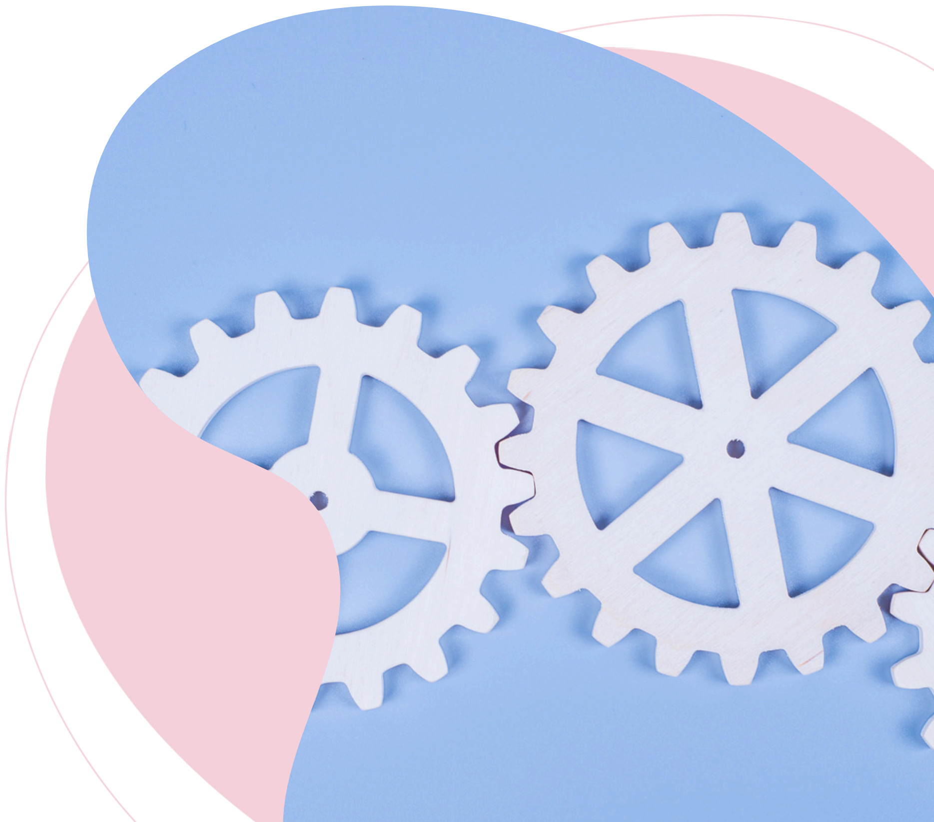

Summary

AI isn't a simple subject, and we've provided a short overview of the basics, which we hope answers some of your questions.

There are many ways in which AI can enhance inventory management and improve warehouse efficiency. The good news is that you don't have to do everything simultaneously. You can introduce AI to your business in stages to make it more affordable.

If you'd like to learn more about how EazyStock uses artificial intelligence, get in touch. We'll be happy to explore how EazyStock can improve your forecast accuracy and product availability while reducing costs.