

AI Glossary of Terms



Simplifying AI.
Empowering Businesses.

Version 1.0 July '25

Table of content

Introduction to AI Glossary..... 2

What's inside.....2

Using This Guide Effectively..... 3

Feedback, Updates, and Version Control.....3

About Navitas AI..... 3

Contact and Next Steps..... 3

Table of content..... 4

Glossary of Terms..... 5

Key Concepts Map: Understanding the Building Blocks of AI & Automation..... 11

 Artificial Intelligence (AI): The Umbrella Term..... 11

 Machine Learning (ML): AI That Learns From Data..... 11

 Natural Language Processing (NLP) & Computer Vision: How AI Perceives the World..... 12

 Generative AI: AI That Creates.....12

 Automation: Putting Intelligence to Work..... 12

 Data: The Fuel for AI..... 13

FAQ's & Common Misconceptions..... 14

Closing Note.....16

Introduction to AI Glossary

As artificial intelligence and automation become more deeply embedded in the way modern businesses operate, the terminology can quickly become overwhelming, especially for those not steeped in the tech world. At Navitas AI, we believe that understanding is the first step toward meaningful adoption.

This glossary has been created to provide clear, straightforward definitions of key terms in the world of AI and automation. Whether you're a business leader exploring AI opportunities, a team member involved in digital transformation, or simply curious about the language of modern technology, this guide is designed to demystify the jargon and help you speak AI with confidence.

What's inside

This resource has been designed to serve as a practical and accessible reference point for professionals navigating the evolving landscape of artificial intelligence and automation.

It begins with a comprehensive **A–Z glossary** of key terms, offering concise definitions written with clarity in mind. These entries aim to provide a shared understanding of the terminology that underpins AI conversations across strategic, operational, and technical levels.

To support broader context, we've included a **Key Concepts Map** outlining the relationships between core ideas such as AI, machine learning, deep learning, and related subfields. This visual framework is particularly useful for those seeking to align AI discussions with business strategy or digital transformation initiatives.

Finally, we address a selection of **Frequently Asked Questions and Common Misconceptions**. This section tackles recurring points of confusion and provides brief, informative responses that clarify the distinctions between overlapping concepts; such as AI versus automation and the practical implications of technologies like generative AI.

Together, these components form a cohesive guide to help business leaders and teams engage with AI from a position of understanding, not uncertainty.

Using This Guide Effectively

To maximise value, we recommend approaching the glossary as a living companion rather than a one-off read. Keep it close during planning sessions, supplier meetings, or board presentations where AI terminology may surface. Each definition has been crafted to stand alone, so you can dip in, locate a term quickly, and return to the task at hand without losing momentum.

Feedback, Updates, and Version Control

Artificial intelligence evolves at pace; so does the language around it. Navitas AI reviews and updates this document on a quarterly cycle, adding new terms and refining existing ones to reflect industry shifts and regulatory developments. The version number and revision date appear on the cover page for clear reference.

Should you spot a gap or require clarity on a specific concept, please contact our team using the details provided below. Your insights help us maintain the glossary's relevance and rigour.

About Navitas AI

Navitas AI is a consultancy dedicated to guiding organisations through the complexities of intelligent technologies. We combine strategic advisory with hands-on solution design to ensure AI delivers measurable business outcomes. Our services span opportunity discovery, automation road-mapping, ethics auditing, and bespoke tool development.

Contact and Next Steps

If you would like tailored support translating these definitions into actionable initiatives, or if your organisation is planning an AI or automation project, we would be pleased to discuss how we can assist.

Email: info@navitasai.co.uk

Web: www.navitasai.co.uk

Glossary of Terms

AI (Artificial Intelligence):

In simple terms, AI refers to computer systems that can perform tasks normally requiring human intelligence. This includes things like learning, problem-solving, understanding language, and recognising patterns. In business, AI is increasingly used for decision support, automation, customer service, data analysis, and product personalisation.

AI Strategy:

A structured approach that defines how an organisation plans to use AI to meet its goals. This includes identifying high-impact use cases, assessing data readiness, choosing tools or partners, and aligning AI initiatives with broader digital transformation plans. A clear AI strategy helps prioritise investment and reduce risk.

Algorithm:

A set of step-by-step rules or instructions that a computer follows to solve a problem or complete a task. Think of it as a recipe for a computer programme. Algorithms underpin everything from search engines and recommendation systems to fraud detection and logistics routing.

Augmented Intelligence:

A model of AI that focuses on enhancing rather than replacing human decision-making. In a business context, this could involve AI systems that support analysts, marketers, or executives by surfacing insights, generating recommendations, or automating parts of their workflow while leaving the final decision to the human.

Automation:

The use of technology to perform tasks or processes with minimal or no human intervention. This can range from simple, repetitive tasks (like sending automated emails) to complex, decision-making processes (such as credit scoring or inventory management). Automation is a core component of digital transformation strategies across industries.

Big Data:

Extremely large and complex sets of information that traditional data processing applications struggle to handle. AI often relies on Big Data to "learn" and identify patterns. Big Data is typically characterised by high volume, velocity, and variety

and is used to power everything from customer segmentation to operational forecasting.

Business Intelligence (BI):

The technologies and practices used to collect, integrate, analyse, and present business data to support better decision-making. AI is increasingly embedded into BI platforms to automate data discovery, generate insights in real time, and power predictive analytics dashboards.

Chatbot:

A computer programme designed to simulate human conversation, either through text or voice. Chatbots are commonly used for customer service or information provision, available 24/7 to handle queries, book appointments, or guide users through processes. Advanced chatbots use natural language processing to better understand user intent.

Computer Vision:

A field of AI that enables machines to interpret and make decisions based on visual data such as images or videos. It's used in applications like facial recognition, quality inspection in manufacturing, and automated checkout systems. Computer vision combines deep learning with image processing techniques to "see" and analyse the world.

Data Governance:

The policies, standards, and processes that ensure high-quality, secure, and ethically managed data across an organisation. Effective data governance is critical to the success of any AI initiative, as poor data quality or unclear ownership can undermine trust in outcomes.

Data Labelling:

The process of tagging or annotating data (e.g. images, text, audio) with relevant information or categories. Labelled data is essential for training supervised machine learning models, ensuring they learn accurately. For instance, labelling thousands of images as "cat" or "dog" helps an AI distinguish between the two.

Deep Learning:

A more advanced type of Machine Learning that uses complex "neural networks" (inspired by the human brain) to learn from vast amounts of data. It's particularly good at tasks like image and speech recognition, language translation, and content generation. Deep learning powers many leading AI applications, including facial recognition and self-driving cars.

Digital Twin:

A virtual representation of a real-world process, system, or asset, often enhanced with AI and real-time data. Businesses use digital twins for simulation, forecasting, and optimisation. For example, modelling a factory floor to improve efficiency or predict maintenance needs.

Ethics in AI:

A growing discipline that considers the moral implications of AI development and use. Topics include transparency, fairness, accountability, data privacy, and avoiding bias in algorithms. Ethical AI aims to ensure that technologies are designed and deployed responsibly, especially when they impact people's lives.

Explainability (Explainable AI / XAI):

The ability to understand and explain how an AI model arrives at its decisions. In highly regulated sectors (e.g. finance, healthcare), explainability is essential for transparency, compliance, and trust. Businesses increasingly require AI outputs to be interpretable by humans, not just accurate.

Foundation Model:

A large AI model trained on a vast dataset and capable of being adapted to a wide range of tasks. GPT-4 is an example. Foundation models serve as a base layer that can be fine-tuned for specific business needs, reducing the cost and time to deploy AI solutions.

Generative AI:

A type of AI that can create new content, such as text, images, music, or even code, rather than just analysing existing data. ChatGPT is a well-known example of generative AI for text. Businesses use generative AI for automating content creation, summarising documents, writing code, and generating design mock-ups, among other applications.

Human-in-the-Loop (HITL):

A process where humans are intentionally kept in the decision-making loop when using AI systems. This ensures oversight, accountability, and quality control. HITL is common in industries like healthcare, finance, and legal services where full automation carries risk.

Intelligent Automation:

The combination of AI technologies (like machine learning and NLP) with traditional automation (like RPA) to enable more adaptive, context-aware workflows. This allows businesses to go beyond rule-based processes and begin automating tasks that involve unstructured data or decision-making.

Knowledge Graph:

A structured representation of facts and relationships between concepts, often used to help AI systems better understand context. For example, a knowledge graph can help an AI know that "Paris" is a city in France, but also that it could refer to a person's name.

Machine Learning (ML):

A branch of AI where computer systems learn from data without being explicitly programmed. Instead of being given specific instructions for every scenario, they learn to identify patterns and make predictions or decisions based on the data they've processed. ML is widely used in fraud detection, recommendation engines, and predictive maintenance.

Model Drift:

Over time, AI models can become less accurate as data patterns change, this is known as model drift. Businesses must monitor and retrain models periodically to ensure performance remains aligned with real-world conditions, especially in dynamic environments like finance or retail.

Model Training:

The process of teaching an AI system to recognise patterns by feeding it data. During training, the model adjusts its internal parameters to improve its accuracy. This is how AI "learns" to classify images, translate languages, or make predictions.

Natural Language Processing (NLP):

A field of AI that focuses on enabling computers to understand, interpret, and generate human language. This is what allows systems to understand what you type or say. From voice assistants and email classifiers to sentiment analysis and language translation tools. NLP is fundamental to chatbots, search engines, and document automation.

Operational AI:

The integration of AI models into day-to-day business processes and systems. While pilot projects and proofs of concept are common, real value is only realised when AI becomes embedded in operations such as automated customer segmentation, fraud detection, or predictive inventory management.

Predictive Analytics:

The use of data, statistical algorithms, and machine learning techniques to identify the likelihood of future outcomes based on historical data. Businesses use this to forecast trends, predict customer behaviour, optimise marketing campaigns, and inform strategic decisions.

Process Automation:

The use of technology to automate business processes, often those that are repetitive and rule-based. This can significantly improve efficiency, reduce operational costs, and free up human teams to focus on higher-value tasks. It's often the first step organisations take in digital transformation initiatives.

Prompt Engineering:

The practice of crafting effective inputs (or "prompts") to get the best possible outputs from generative AI tools like ChatGPT or image generators. As these models don't follow traditional code logic, the way you ask a question can significantly influence the quality and accuracy of the response.

Proof of Concept (PoC):

A small-scale project designed to test whether a specific AI or automation use case is technically and commercially viable. Businesses use PoCs to reduce risk before scaling up investment, typically measuring feasibility, performance, and value.

Reinforcement Learning:

A type of machine learning where an AI learns by trial and error, receiving rewards or penalties for its actions. This method is used in areas like robotics, game playing (e.g. AlphaGo), and recommendation engines, where the system improves by learning from outcomes over time.

Responsible AI:

A framework for ensuring AI technologies are developed and deployed in ways that are ethical, transparent, and aligned with societal values. For businesses, this includes practices like bias testing, explainability, consent management, and data security especially important when customer or employee data is involved.

Return on AI (ROAI):

A business metric that quantifies the value derived from AI initiatives relative to the investment made. It considers both tangible benefits (e.g. cost savings, revenue growth) and intangible ones (e.g. faster decisions, customer experience improvements). Tracking ROAI helps justify and prioritise future AI spend.

Robotic Process Automation (RPA):

A specific type of process automation that uses "software robots" (bots) to mimic human actions when interacting with digital systems and software. RPA bots can open applications, enter data, extract information from documents, and complete routine tasks just like a person would but with greater speed and accuracy.

Supervised Learning:

A type of Machine Learning where the AI is trained on data that has already been

labelled. For example, showing it many pictures of cats and dogs, with each picture clearly marked as "cat" or "dog," so it learns to identify them. This is commonly used in applications like spam detection, image recognition, and medical diagnostics.

Synthetic Data:

Artificially generated data that mimics real-world data. It's used to train or test AI models when actual data is scarce, sensitive, or too costly to collect. For example, synthetic patient records might be used to train a healthcare AI model without compromising privacy.

Training Data:

The data used to teach an AI or machine learning model. The quality, quantity, and diversity of training data have a direct impact on model performance. Biased or incomplete training data can lead to flawed results.

Transfer Learning:

A technique where a model trained on one task is adapted to perform another, related task with minimal additional training. This allows organisations to use pre-trained models as a starting point, saving time and resources while maintaining high performance.

Unsupervised Learning:

A type of Machine Learning where the AI is given unlabelled data and has to find patterns and structures within it on its own. This is often used for tasks like grouping similar customers together (clustering), anomaly detection, or exploring data to find hidden relationships.

Key Concepts Map: Understanding the Building Blocks of AI & Automation

Artificial intelligence can often feel abstract, a tangle of overlapping terms, emerging technologies, and unclear boundaries. To help make sense of it all, we've included a **Key Concepts Map** that visually outlines the core ideas behind AI and automation, and how they relate to one another in practice.

This framework is designed to provide clarity for decision-makers, teams, and stakeholders looking to align technical understanding with strategic objectives.

Here's a breakdown of how these concepts connect:

Artificial Intelligence (AI): The Umbrella Term

At the highest level, AI refers to systems that mimic human intelligence. Under this broad category sit the specific disciplines and tools businesses interact with most.

Machine Learning (ML): AI That Learns From Data

Machine Learning is a core subset of AI. Rather than being programmed with fixed rules, ML algorithms learn patterns from historical data. It powers everything from fraud detection to product recommendations.

- **Supervised Learning** uses labelled data to teach the model.
- **Unsupervised Learning** finds hidden patterns in unlabelled data.
- **Reinforcement Learning** improves performance through trial and error.
- **Deep Learning** is a more complex form, using layered neural networks.

Natural Language Processing (NLP) & Computer Vision: How AI Perceives the World

Two major application areas of ML:

- **NLP** allows AI to understand, interpret, and generate human language. It's what powers chatbots, sentiment analysis, and document summarisation.
- **Computer Vision** enables machines to interpret images or video. This is useful in areas like quality control, facial recognition, and scanning documents.

Generative AI: AI That Creates

Generative AI refers to systems that can produce new content such as text, images, code, or even music. Models like ChatGPT and DALL·E fall into this category. It's transforming content creation, marketing, product design, and more.

- Requires a foundation of machine learning, typically deep learning.
- Often built on **Foundation Models** that are then adapted for specific tasks.

Automation: Putting Intelligence to Work

Automation applies AI (and other technologies) to streamline processes.

- **Process Automation** refers to any technology automating repetitive tasks.
- **Robotic Process Automation (RPA)** mimics human actions in digital systems.
- **Intelligent Automation** combines RPA with AI capabilities like NLP or ML for smarter, context-aware workflows.

Data: The Fuel for AI

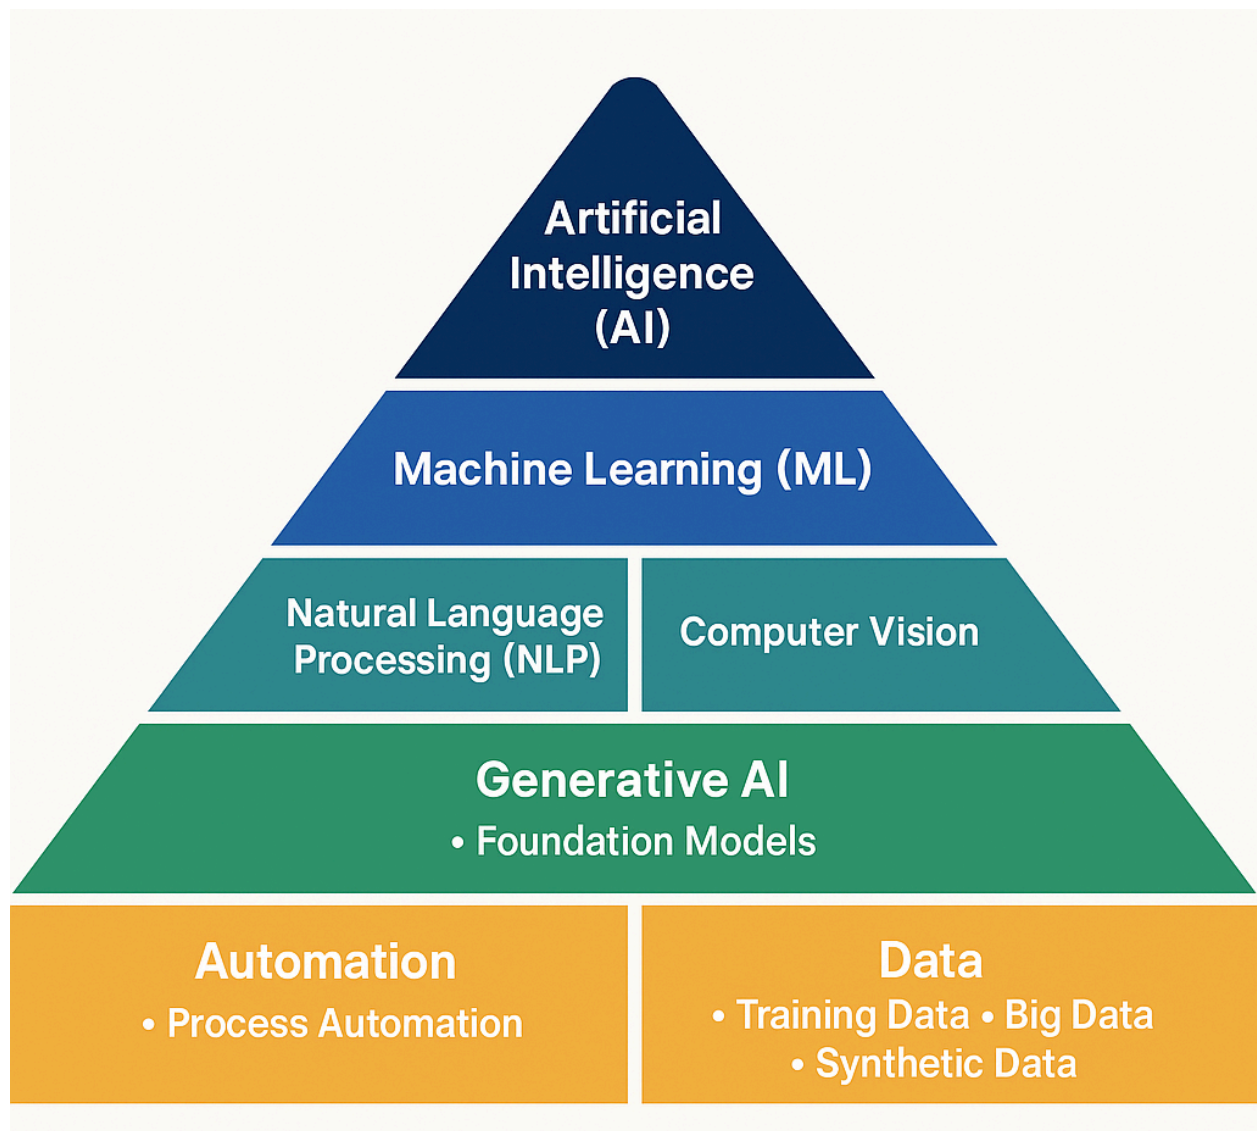
All of the above relies on quality data. Businesses must have:

- **Training Data** to build models
- **Big Data** for scale and pattern recognition
- **Synthetic Data** where real data is unavailable or sensitive
- **Data Labelling** and **Data Governance** to ensure accuracy, security, and compliance

This map is not a static diagram. It represents a system that evolves as technologies and use cases progress. Whether you're beginning your AI journey or scaling existing initiatives, understanding how these concepts fit together is key to making strategic, confident decisions.

Key Concepts

Understanding the Building Blocks of AI & Automation



FAQ's & Common Misconceptions

As AI and automation become more prevalent in modern organisations' psyche¹⁷, certain questions and misunderstandings often arise. This section addresses a selection of the most common, offering clarity for business leaders, project teams, and decision-makers alike.

"Is AI the same as automation?"

Not exactly. Automation refers to technology performing tasks with minimal human input, often repetitive or rule-based. AI, on the other hand, involves systems that can learn, adapt, and make decisions based on data. While AI can power more intelligent forms of automation, not all automation involves AI.

"Will AI replace all jobs?"

AI will transform jobs, not eliminate them. While some roles may be automated, many more will be augmented; freeing people from repetitive tasks and creating demand for new skills in oversight, strategy, and AI management. The greatest value comes when AI enhances human capability, not replaces it.

"If I use AI, do I need a data scientist?"

Not necessarily. Many AI tools are now accessible through no-code or low-code platforms, designed for business users. However, for more complex or large-scale deployments, technical expertise in areas like data engineering, machine learning, or governance may be required either in-house or via a partner.

"Is ChatGPT the same as AI?"

ChatGPT is a **type** of AI; specifically, a generative AI language model. While powerful, it represents just one of many applications under the AI umbrella. AI also includes image recognition, predictive analytics, robotic automation, and much more.

“Does AI always give the right answer?”

No. AI systems can be highly accurate but are not infallible. They depend on the quality of training data, the model's design, and the context in which they're used (the same as humans). It's important to validate outputs, monitor for bias or drift, and involve human oversight where appropriate.

“Is more data always better?”

Not always. While large datasets can improve model performance, the quality, relevance, and structure of data are just as important. Clean, well-labelled data aligned with your specific use case will often outperform larger but noisier datasets.

“Can we plug in AI and expect instant results?”

AI is not a plug-and-play solution. Successful implementation requires a clear problem definition, readiness assessment (data, processes, skills), and a practical approach to testing and scaling. The best results come when AI is integrated thoughtfully into workflows and not added as an afterthought.

Closing Note

This glossary has been designed as a practical reference for navigating the language of artificial intelligence and automation in a business context. As the field evolves, so too will the terms and technologies that define it.

If you've found this guide useful and would like support translating these concepts into meaningful action, the team at Navitas AI is here to help. From strategy development and solution design to implementation and governance.

Let's turn understanding into impact.

Contact us:

 info@navitasai.co.uk

 www.navitasai.co.uk



Simplifying AI.
Empowering Businesses.