Welcome to the third Market Scan from the IAASB's Disruptive Technology team. Building on our previous work, which included the Innovation Report created with Founders Intelligence and discussed at the January 2021 IAASB Meeting, we issue a Market Scan focusing on topics from the report approximately every two months. Market Scans consist of exciting trends, including new developments, corporate and start-up innovation, noteworthy investments and what it all might mean for the IAASB.

In this Market Scan, we explore Artificial Intelligence (AI), which is used in a broad range of technologies across the audit and assurance value chain. This Market Scan provides a high-level primer on Artificial Intelligence as it is one of the most significant and potentially disruptive technologies in audit and assurance. Future Market Scans will build on this by focusing on some of the specific AI-powered technologies highlighted below.
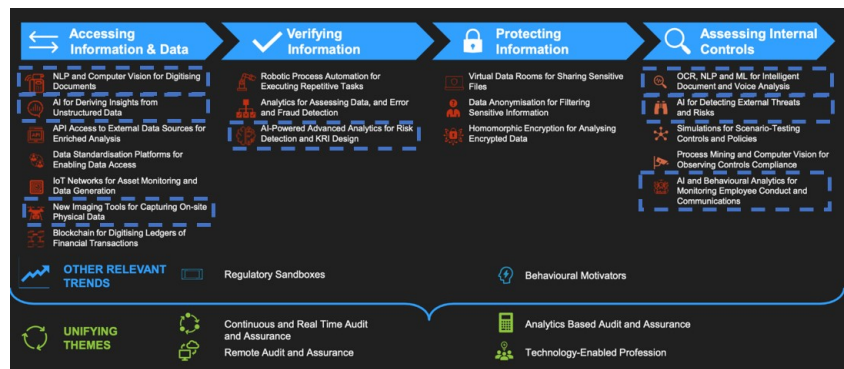
See full graphic in the Innovation Report, available on the IAASB website.

We cover:

- What is AI, including related concepts of machine learning and deep learning
- AI use cases in audit and assurance
- AI challenges
- AI developments

What Is Artificial Intelligence?

Artificial Intelligence (AI) is a broad discipline of computer science that refers to the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision making, and translation.

AI also describes a broad range of technologies shown in the diagram below. Many of the technologies we use every day contain one or more of these capabilities; for example, a smart speaker contains speech recognition (to turn our speech into text), natural language processing (NLP) (to understand the request and generate a response) and machine learning (to improve the quality of responses over time).

Overview of AI Technologies
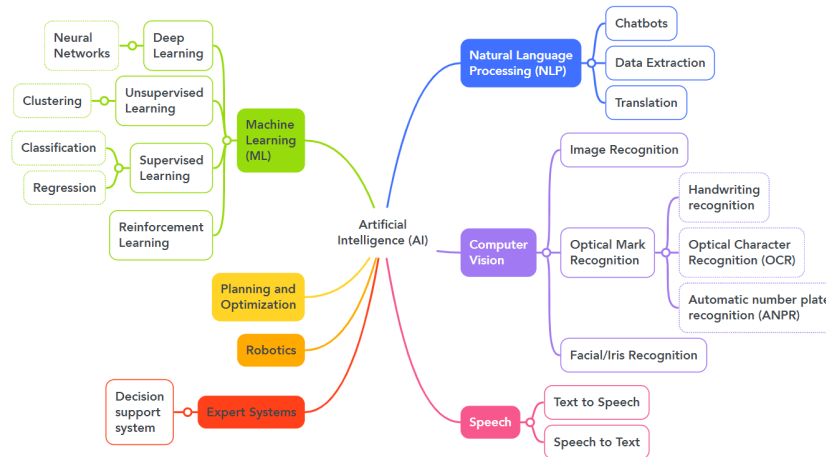
Intelligence in this context is the ability to perceive or deduce information, retain it as knowledge and apply it to making decisions. In computers, this is done by analyzing large quantities of data using advanced statistics (including probability analysis) to find patterns and make predictions.

Types of AI

Narrow AI (Today's AI, weak)	General AI (Future AI, strong)
Applications that model human behavior to perform a specific task or function, e.g., face recognition, speech detection	Currently hypothetical but refers to machines that have full human cognitive abilities

What is an algorithm?

Algorithms are in use all around us, although the term may not be fully understood as frequently. Think of it as a recipe used by computers: a finite sequence of well-defined instructions, typically used to solve a class of specific problems or to perform a computation. An algorithm takes an input (e.g., a dataset) and generates an output (e.g., a pattern that it has found in the data). It is like taking your ingredients and following a recipe to bake a cake.

Algorithms are not exclusive to AI. They are likely used in every audit to complete procedures such as identifying sample sizes or performing data analytics, such as ratio or regression analysis.
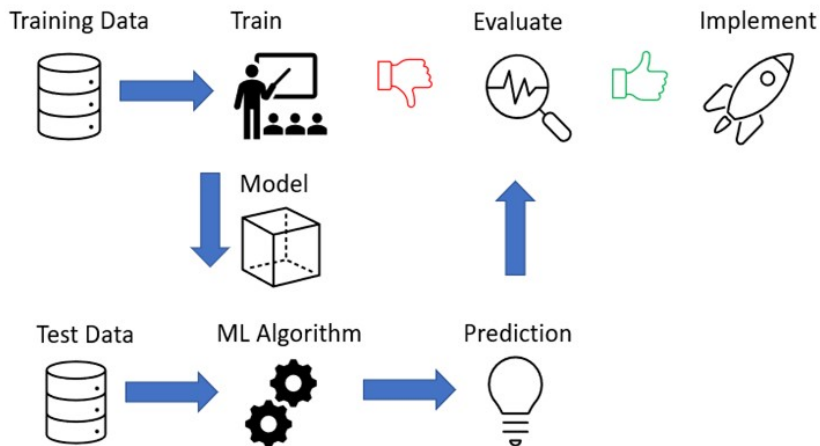
What Is Machine Learning?

Machine learning is about using algorithms to guide predictions. The goal of the machine learning process is to create a model, which is based on one or more algorithms. The model is developed through training with the goal that the model should provide a high degree of predictability.

One of the earliest examples of a machine learning system was a computer checkers game created by Arthur Lee Samuel at IBM. Arthur demonstrated how machine learning could work by creating a computer function to measure the chance of winning based on the position of pieces on the board. The computer then used function to determine the move most likely to lead to a successful outcome, that is, winning. The computer learns by using the feedback from playing as its data and using Arthur's function to guide its prediction model to get to a preferred outcome.

In its simplest form, machine learning requires a five-step process:

1. Get and organize the data
2. Choose a model (one or more algorithms)
3. Train the model (using training data, about 70% of your data set)
4. Evaluate the model (using test data, about 30% of your data set)
5. Fine-tune the model and implement

Machine Learning Process

The main challenges with machine learning implementation are in relation to what data to use (and how to get it) and what model to use, that is, which algorithms to apply.

Machine learning approaches

There are three main types of learning approach used in machine learning; determining which approach to use largely depends on what data you have available.

Supervised learning is an approach used when large amounts of labeled data are available. This enables the technology to learn by comparing its results to the correct answer. There are effectively two types of algorithms that are used within supervised learning—one is classification, where you divide the dataset into common labels. A common form of classification algorithm is called Naïve Bayes Classifier, which is used in text analysis (e.g., for sentiment analysis, email spam detection). It uses frequency and patterns in data to come up with a prediction model based on probabilities.

The other type of algorithm used in supervised learning is regression, which finds continuous patterns in data. A common form of regression algorithm is linear regression, which shows the relationship between variables and uses this to predict outcomes based on inputs, e.g., predicting expected sales per square foot of sales floor space.

Supervised vs Unsupervised Machine learning: What's the difference?

(Eye on Tech video, two-minute watch)

Unsupervised learning is used when the available data is unlabeled, so the algorithms used seek to put the data into groups. The most common approach is called Clustering, which is grouping similar items together and then iterating the model to get better results. There are a variety of quantitative methods, i.e., ways of grouping items. Common uses of unsupervised learning are customer segmentation for targeting marketing messages where similar customer characteristics are expected to share similar preferences.

Finally, **Reinforcement learning** is commonly used in gaming and robotics, effectively learning through a process of trial and error to get the most effective outcome (such as winning the game or navigating successfully around a space).

Useful Resources

- General AI vs Narrow AI | Levity, five-minute read
    - differences and examples
- Artificial intelligence explained in 2 minutes: What exactly is AI? | KI-Campus, two-minute video
    - simple explanation with examples
- AI and Machine Learning Explained with Examples | ClassicInfomatics, seven-minute read
    - machine learning methods explained with examples
- Top 15 Hot Artificial Intelligence Technologies | Edueka!, six-minute read
    - short explanations of key AI technologies

What Is Deep Learning?

Deep learning is a subfield of machine learning that uses neural networks for learning and bear some resemblance to how the human brain works. This way of processing data is more granular than with machine learning and involves more layers of analysis. Although the concept of deep learning has been around since the 1970s, its recent growth is due to the significant advancements in computing power. It is commonly used for speech and image recognition.

An artificial neural network ingests data through an input layer, processes it through a complex network (known as the hidden layer or layers) to provide an output. The word "hidden" in the hidden layer simply refers to the fact that the units in the layer are not visible to external systems and are "private" to the neural network.
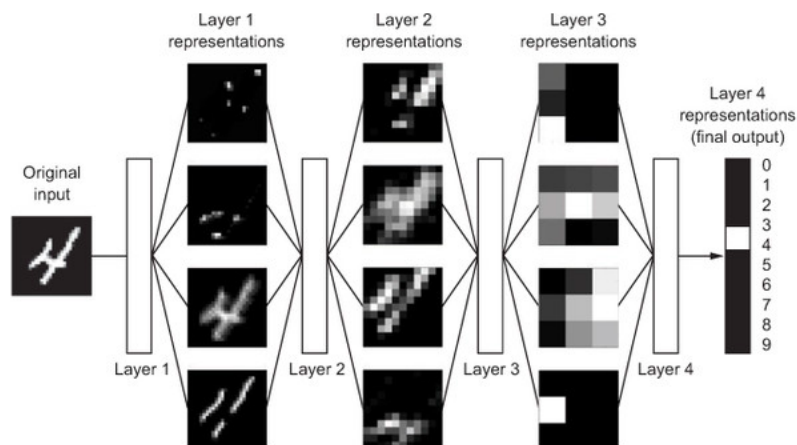
Example of a neural network used to identify the number 4
(From Deep Learning with Python by Francois Chollet)

Each of the processing units in the network is called a neuron. A neuron is a container with an input value, a weighting, and a bias (which is a constant). These are computed together and then an activation function is applied, which is effectively a mathematical operation that normalizes the inputs and produces an output that is then passed onto neurons in the next layer.

The weightings along with the bias can change the way the neural networks operate and are used to refine the model to get to the preferred outcome.

The most common types of neural networks are called fully connected neural networks, referring to all the neurons having connections from layer to layer. Other neural networks include recurrent neural networks, convolutional neural networks and generative adversarial networks.

In **Recurrent Neural Networks (RNNs)**, the function not only processes the input but also prior inputs across time. An example of this is with predictive text, as you start to type, different word options are presented based on what the system predicts you are typing.

In **Convolutional Neural Networks (CNNs)**, data is processed in stages from easy to complex with each of the stages being a convolution. CNNs are often used in computer vision applications such as image recognition software.

Generative Adversarial Networks (GANs) are a relatively new but powerful class of neural network used for unsupervised learning. They are made up of a system of two neural network models (a generator and a discriminator) that compete with each other and are able to analyze, capture and copy the variations in a dataset. It is this technology that gave rise to creation of deepfakes; they have also begun to be used by the financial services sector to help with fraud identification.

Useful Resources

- What is a Neural Network? | DeepAI, four-minute read
  - simple explanation and example
- What is a Neural Network? | Simplilearn, five-minute video
  - neural network basics and examples

AI Use Cases in Audit and Assurance

There are many ways that AI may be deployed to support the audit process.

Audit Planning

- Resource optimization using AI technology to analyze staff profiles and experience to bring together the best team for the type of audit engagement
- Client acceptance procedures using AI to analyze data from non-traditional sources, such as social media, emails, phone calls, public statements from entity management, etc., to identify potential risks relevant to client acceptance and continuance assessments.

Understanding the entity and its systems, and identifying risks

- Using natural language processing and machine learning AI technologies to analyze structured and unstructured information, such as global regulatory notices, industry reports, regulatory penalties, news, public forums, etc., to detect risks of audit relevance
- Intelligent document analysis, such as optical character recognition natural language processing and machine learning technologies, to derive insight from unstructured data sources like email, documents, transcribed voice, images, etc. to support understanding of the entity's information system and related controls.
- Quickly and more efficiently understanding the entity's internal controls by summarizing and extracting what has been documented in process documents, emails, articles, and from employee inquiries.
- AI-powered behavioral analytics to identify suspicious or unusual entity employee behavior and intent, such as data exfiltration, employee collusion or abuse from privileged users.
- Enhancing an audit team's judgments on higher-risk areas of audit engagements by using AI to identify common risks relevant to the entity's industry, regulatory environment, operating locations and other external factors.

Substantive Procedures

- AI tools, benefiting from increases in the quality and quantity of available "training" data, can be applied to data sets to algorithmically identify outliers and anomalous data and to perform predictive analytics for use in areas such as testing large transaction populations, auditing accounting estimates and going concern assessments.
- Document processing, review and analysis by using optical character recognition to identify and extract key details from contracts (e.g., leases) and other documents (e.g., invoices)
- Inventory and physical asset verification procedures through use of drones with computer vision (image recognition) particularly for larger capital assets, such as trucks, or the inspection of large-scale business sites, such as wind farms.

Conclusion Procedures

- AI technologies to support auditors' work on financial statement disclosures enabling easier identification of missing disclosure requirements and non-compliance.
- AI technologies to support tick and tie of underlying audit work through to financial statements and related disclosures

Some of these technologies will be explored in more detail in future Market Scans.
Many organizations are expanding their use of AI across parts of their business with the goal of driving operational efficiencies, better informed decision making and generating growth through innovation. As a result, it is likely that this technology will become a relevant consideration when performing audit procedures, particularly regarding risk identification and assessment, and risk response activities.

Useful Resources

- The Intelligent Audit | ISACA, nine-minute read
  - examples of use of AI in the audit
- AI in the Accounting Big Four – Comparing Deloitte, PwC, KPMG, and EY | Emerj, 11 minute read

AI Challenges

Where AI is deployed, whether by the auditor in carrying out their procedures or by an audited entity within their business operations, the associated risks need to be identified and appropriately managed. Many assurance firms and organizations have developed methodologies that provide a framework for identifying and managing AI related risks. In September 2021, COSO issued new guidance setting out how to apply "the COSO framework and principles to help implement and scale artificial intelligence". This guidance identifies five areas of AI related risks:

- Bias and reliability breakdowns due to inappropriate or non-representative data
- Inability to understand or explain AI model outputs
- Inappropriate use of data
- Vulnerabilities to adversarial attack to obtain data or otherwise manipulate the AI model
- Societal stresses due to rapid application and transformation of AI technologies

It concludes that appropriate risk management is needed to ensure that AI solutions are "trusted, tried and true".

Auditing AI may require a different set of skills to those currently applied in today's audits and many firms are updating their recruitment strategies, training curricula and audit methodologies to respond to the growing need for AI competencies. Future Market Scans will explore some of these challenges in more detail.

Useful Resources

- 6 lessons from audit experts who adopted AI early | Journal of Accountancy, five-minute read
  - transcript from interview with auditors who have implemented AI

- Auditing Artificial Intelligence | Corporate Compliance Insights, six-minute read
  - impact of AI on planning an Internal audit
- COSO Releases New Guidance: Realize the Full Potential of Artificial Intelligence | COSO press release, two-minute read

Audit and Assurance Publications

- The Data-Driven Audit: How Automation and AI are Changing the Audit and the Role of the Auditor | AICPA
- A CPA's introduction to AI: From algorithms to deep learning | CPA Canada

---

AI Developments

The global AI market is expected to achieve a compound annual growth rate of nearly 40% over the next five years and whilst AI technologies such as natural language processing and speech recognition are maturing, others such as deep learning and Generative AI have significant scope for development.

Here are some recent noteworthy developments:

Regulation and Explainable AI
One of the issues that has arisen with AI is around the negative impact of biases in algorithms and the harm that this can cause. In a recent survey more than one in three companies surveyed disclosed that they had suffered losses (revenue, customers or staff) due to AI bias in their algorithms. In response, there is an expectation regulation will be established in the near future. The EU, in its white paper, "On Artificial Intelligence—A European Approach to Excellence and Trust", noted that explainability is a key factor to improving trust in AI. Many companies are, therefore, expected to look to implement explainable AI in which the results of the solution can be understood by humans.

Efficient AI
DeepMind, the company behind the AlphaGo program that was the first to beat a professional Go player, has developed an AI large language model—that is, a statistical tool to predict words—called RETRO (Retrieval-Enhanced Transformer). This AI technology, built to generate convincing text, chat with humans and answer questions is said to match the performance of neural networks 25 times its size through use of a text database.

Decision Intelligence
One of the top technology trends for 2022 noted by Gartner is decision intelligence, which is using AI to enhance and support human decision making. Peak.ai, a UK-based start-up, raised US $75m in series C funding in August 2021 to enable it to build out its "decision intelligence" platform to expand into new markets and help non-tech companies make AI-based decisions.

Funny Story
AI argues for and against itself in Oxford Union debate: Megatron, an AI developed by Google and Nvidia, was given access to huge quantities of data to enable it to both defend and argue against the motion, "This house believes that AI will never be ethical". It's not clear which argument was more compelling!

What do you think about this bulletin?

Please take the time to fill out our quick survey to let us know your thoughts about this bulletin, how it can be improved and what you would like to hear about going forward.

Take our survey