

| Term | Definition |
|---|---|
| Accuracy | The most common performance measure for Machine Learning. It is an expression of how often the AI system provides a correct response, as opposed to a false positive or false negative. |
| Adversarial Examples | A security threat unique to machine learning. It is the intentional inclusion of an example to mis-train an AI. Adversarial Examples would cause AI to develop a vulnerability, for example an Autonomous Vehicle could be trained to not recognize a stop sign. |
| Algorithms | The "gears" that make Artificial Intelligence work. A set of instructions or mathematical functions that deliver a command or output from interpretation and analysis of inputted data. Algorithms can perform calculations, data processing and analysis, classification, and reasoning, to name a few. |
| Application Programming Interface (API) | Application Programming Interface (API): Comprises a set of definitions, protocols and tools for building application software. Simply put, it is a collection of methods that allow software components to interact with each other, allowing, for example, users to copy and paste text or other types of data from one application to another. There are many different types of APIs for operating systems, applications or websites. A good API facilitates the development of a programme by providing all the building blocks, which are then assembled by a Programmer. |
| Artificial General Intelligence (AGI) | Artificial general intelligence (AGI): Refers to the idea that general human intelligence – which spans different domains and capabilities – could be matched or even surpassed by machines. Some research uses the term "Artificial Super Intelligence", or "superintelligence", to refer to hypothetical General AI systems that could far surpass the abilities of humans. |
| Artificial Intelligence (AI) | Artificial Intelligence (AI): No universally accepted definition for AI exists. AI is an umbrella term initially coined in the 1950s to talk about systems, machines or computers who can seem to imitate how humans think. Intelligent technology, programs and the use of advanced computing algorithms that can augment decision making by identifying meaningful patterns in data. A blanket term which describes technologies that allow computers and machines to perform tasks that require human intelligence or capabilities like visual perception, decision making, reasoning, and translation. Today, the term can be used to describe different things: 1. AI as a field of research that covers the design of intelligent systems, the development of associated methods and techniques but also includes many other considerations such as the ethical and societal impact of such systems. 2. AI as a technology (i.e. the application of this body of knowledge for solving real-life problems). Examples of these technologies include machine learning and neural networks. AI can scour massive amounts of data to detect patterns in a short amount of time. |

| | |
|---|--|
| <p>Artificial Intelligence Augmentation (Intelligence Augmentation)</p> | <p>Artificial Intelligence augmentation (intelligence augmentation): Relates to instances where humans and machines learn from each other and redefine the breadth and depth of what they do together.</p> |
| <p>Artificial Intelligence Effect (AI Effect)</p> | <p>Artificial Intelligence effect (AI effect): The phenomenon whereby "as soon as Artificial Intelligence (AI) researchers achieve a milestone long thought to signify the achievement of true artificial intelligence, e.g. beating a human at chess, it suddenly gets downgraded to not true AI."</p> |
| <p>Artificial Intelligence System</p> | <p>Artificial Intelligence system: A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy. In addition, AI are "machines performing human-like cognitive functions".</p> |
| <p>Artificial Narrow Intelligence (ANI)</p> | <p>Artificial narrow intelligence (ANI): A term used to recognise that all current AI applications can be used effectively to perform.</p> |
| <p>Artificial Neural Networks (ANNs)</p> | <p>A collection of connected artificial neurons designed to perform calculations to complete specific tasks. The simplest form of ANNs is composed of three layers of neurons: an input layer, a hidden layer and an output layer. All neurons from the input layer are connected to all neurons from the hidden layer, which are connected to all neurons in the output layer.</p> |
| <p>Artificial Neuron (or Node)</p> | <p>A mathematical model inspired by the biology of human brains. It usually has many inputs and a single output linked by a mathematical function, also called the activation function or transfer function. Several neurons connected to each other can be used to create an artificial neural network.</p> |
| <p>Augmented Intelligence</p> | <p>The current state of Artificial Intelligence. Human ability has been amplified by the use of these technologies. Pattern recognition and information processing are outsourced to computers and made much quicker than could ever be done by humans. Humans can review recommendations or analysis by computers to aid in decisions or other work. Instead of replacing humans, these technologies compliment their abilities by making them faster. However, the machines are still not "intelligent" in themselves in needing to be directed in their learning and their generally narrow applications.</p> |
| <p>Automated Planning</p> | <p>Refers to the ability of machines to automatically and autonomously devise courses of action or strategies to achieve a goal, including anticipating the effects of different approaches.</p> |

| | |
|-----------------|--|
| Bias | Bias in data refers to the systematic distortion in the sampled data that compromises its representativeness, in algorithms it describes systematic and repeatable errors in a computer system that create unfair outcomes, such as privileging one arbitrary group of users over others. |
| Blockchain | An emerging system for digital storage that relies on decentralization and encryption to make stored data more resistant to attack and keep records safe. Some have advocated developing blockchain alongside AI to keep sensitive data more secure, help researchers follow algorithmic thought processes (which are currently very difficult or impossible to ascertain), and because AI will be able to manage blockchains better than humans |
| Chatbot | AI technologies that are able to communicate with humans, usually through computer text. Commonly, they are deployed as customer service reps to handle consumer complaints and either offers a simple answer, assist in navigation of a website, or directing the customer to a proper representative. |
| Classification | Process of categorizing a group of objects while only using some data features about objects. It is a common application of AI. For example, identifying a picture as containing a dog or not, identifying applications that should pass or be rejected, and flagging outliers in datasets like tax records to detect fraud. |
| Computer Vision | Refers to the ability of AI to process and synthesise visual data (e.g. detect and classify objects based on images or videos) and perform tasks such as facial recognition and scene interpretation. In addition to processing existing visual data, some applications involve the construction of visual data, such as the creation of 3D models from 2D images. |
| Data Feature | An explanatory variable in a dataset that is used to make predictions |
| Data Fusion | Merging of multiple heterogeneous datasets into one homogenous representation to be better processed for data mining and management. |
| Data Governance | Refers to a system of decision rights and accountabilities for information-related processes, executed according to agreed-upon models which describe who can take what actions with what information, and when, under what circumstances and using what methods. |
| Data Labelling | Involves identifying elements that an AI system will try to predict. It may also be referred to as the answer, result or output. |
| Data Lifecycle | The entire period of time that data exists in your system. This life cycle encompasses all the stages that your data goes through, from first capture onward. |

| | |
|----------------------|--|
| Data Mining | Analysis of large datasets to look for relationships between variables. Often this is done in an automated way where a computer does a rapid succession of tests to identify relationships. Humans must judge the substantive value and validity of those relationships. |
| Data Normalisation | The conversion of values from their raw state into a standard range of values. |
| Data Quality | Term used to describe a documented agreement on the representation, format, and definition for data. |
| Data Silos | The phenomena where different government agencies have troves of data that is isolated within that agency, disconnected from the data of other agencies. If the data was more integrated and connected, more insights, and more valuable algorithms could be developed. |
| Data use sensitivity | Risks or considerations associated with data subjects themselves or use of data. |
| Deep Learning | A form of Machine Learning for which the design is inspired by the biology of human brains. Deep Learning works by exposing multi-layered artificial neural networks (see above) to vast amounts of data. Any artificial neural network that learn a long chain of causal links; more complicated neural networks with more nodes ("neurons") and connections between them. AI with more advanced neural networks can go beyond learning specific tasks and also identify the features of the task. This creates a better understanding of what is being done by the algorithm, and may allow it to generalize its task or more effectively improve its performance. It is described by some as the AI learning to ask "why". For example, instead of simply learning to identify pictures with dogs, deep learning will identify fur, paws, and eyes as separate components of a dog. |
| Expert Systems | An application of AI. The AI is programmed to make decisions that typically require human level expertise in things like making medical diagnosis or driving a car. These technologies are largely behind automation. |
| Feature Engineering | The process of determining which features might be useful in training a model. |
| Features | Any identifying characteristic AI uses to analyse, compare, or make decisions. In visa applications this may be a place of origin or criminal history; word choice for natural language processing; or eye colour in facial identification. |

| | |
|--------------------------------|--|
| <p>General Intelligence</p> | <p>Human level intelligence which is able to do things like make connections between tasks, combine skills, become cognizant, and other more complex abilities. The dream of AI would be to develop a program that could reach this level of cognition. Currently, AI systems are unable to operate outside of their narrowly defined tasks and are not considered "Generally Intelligent"</p> |
| <p>Generalisation</p> | <p>Refers to an AI model's ability to make correct predictions on new, previously unseen data as opposed to the data used to train the model.</p> |
| <p>Harm</p> | <p>Any adverse effects experienced by an individual (or organisation) including those which are socially, physically, or financially damaging.</p> |
| <p>Human Rights</p> | <p>Are rights inherent to all human beings, regardless of race, sex, nationality, ethnicity, language, religion, or any other status. Human rights include the right to life and liberty, freedom from slavery and torture, freedom of opinion and expression, the right to work and education, and many more. Everyone is entitled to these rights, without discrimination.</p> |
| <p>Inputs</p> | <p>AI requires a way to perceive its spatial or analytic environment, depending on the functions; this is the data it uses. Visual, audible, numerical, or qualitative data can be the input and it can be structured or unstructured. But not all AI can take all data, it depends on what it has been programmed to analyse.</p> |
| <p>Knowledge-Based Systems</p> | <p>Capture and store facts in a "knowledge base", and then use an "inference engine" to infer insights from the knowledge base in order to solve problems, often through programmed IF-THEN rules.¹⁹² For example, "expert systems" codify the knowledge of experts to help address problems through IF-THEN rules.</p> |
| <p>Machine Learning</p> | <p>The process a machine goes through to teach itself a task and improve its completion of the task. Data is input, analysed, and an output is returned. The machine can then compare its performance with a training baseline or its standalone performance (supervised or unsupervised learning) and improve its performance without human intervention. But, there will likely be human oversight to evaluate its performance. In other words, the machine creates a (statistical) model of what it is doing or investigating and refines that model as more information is given to it. For example, the model may learn that fur is not a strong identifying characteristic of dogs after it wrongly uses fur to classify a cat as a dog. A subset of AI and an approach where machines learn to make predictions in new situations based on historical data. Machine Learning consists of a set of techniques to allow machines to learn in an automated manner, without explicit instructions from a human, by relying on patterns and inferences. Machine Learning approaches often teach machines to reach an outcome by showing them many examples of correct outcomes – called "training". Another approach is for humans to define a set of broad rules and generally let the machine learn on its own by trial and error.</p> |

| | |
|-----------------------------------|---|
| Machine Learning Generalisation | A phase after validation in which a machine learning system is deployed in a real environment. |
| Machine Learning Training | A phase in which an AI system is exposed to data which it learns from by applying statistical models. |
| Natural Language Processing (NLP) | Refers to computers' ability to handle and interpret human language and perform various tasks such as translation or text analysis. Computer translating and understanding of human language. For example, your iPhone personal assistant hearing your voice, understanding what you requested, and returning results or performing voice to text functions. Other examples include translators and chatbots. |
| Nearest Neighbour | Machine Learning where the algorithm compares its performance to the most similar observation, its 'nearest neighbour', or nearest Kth neighbours depending on the researchers' desires. Typically, this works in the realm of classification, like Amazon's recommendations or ad services. To recommend a new movie, Netflix will identify the 'nearest neighbours' to the movie you just rated 5 stars. |
| Neural Networks | A specific technology of AI that tries to simulate communicating brain neurons and allows it to learn. Input is split into separate simulated neurons which given weights for their importance in analysis. The neurons perform a process and "score" their data. In other terms this means the input is analysed, and the information is sent to other neurons for other analysis until it reaches the output, where a decision is made. The adjustment of weights given to each node is part of the learning process. Nodes must determine where to pass along the information they are given, which comprises how they are connected within the neural network. Output nodes may be classifiers, like accept or reject. In this case, the node with the higher score is the deciding one. The middle notes would identify if the information matches what it knows about passing or not passing observations and scores based on that. |
| Non-Operational AI | Systems that do not use a live environment for their source data. Most frequently, they produce analysis and insight from historical data. |
| Operational AI | Systems that have a real-world effect. The purpose is to generate an action, either prompting a human to act, or the system acting by itself. Operational AI systems often work in real time (or near real time) using a live environment for their source data. |
| Outputs | AI output can be broad. From real-time traffic data the output may be a red or green light on a traffic signal. In visa applications, the AI can return recommendations like accept or reject. The AI can return correlations or patterns in data to flag outliers or provide new information. It may also be a response to a customer question. The output that AI can perform is diverse. |

| | |
|----------------------------------|--|
| Overfitting | Refers to cases when an algorithm is too specific to the extent that it captures and focuses too much on noise and anomalies. |
| Propagation | The process of working through a neural network. There is forward propagation, or moving from the inputs to the outputs. There is also backward propagation, where the machine works backwards and assesses its process and can adjust the weights or making other changes to refine itself, this is how it learns or trains itself. |
| Reinforcement Learning | Machine learning where the algorithm checks its performance against a "correct" example given by a human. A form of ML that works by having an agent (computer) complete a task by interacting with an environment. Based on these interactions, the environment will provide feedback that causes the agent to adapt its behaviour. In other terms, the agent learns through trial and error, where error is penalised by the environment and success rewarded. It then automatically adjusts its behaviour over time producing more refined actions. It can identify if it was right or wrong and adjust based on that outcome. It can also identify numerous solutions for humans to review. Researchers can set rewards or punishments for behaviour can be set to direct the development of the algorithm, and the algorithm will then seek to optimize itself by minimizing punishment and maximizing reward. One must be careful in setting these constraints as they will be taken very literally and precisely by the computer. |
| Responsible Officer | The Officer who is responsible for: use of the AI insights / decisions; the outcomes from the project; the technical performance of the AI system; data governance. |
| Reversible Harm | Refers to an adverse effect that can be reversed with some level of effort, cost and time. |
| Robotic Process Automation (RPA) | A business process automation technology that automates manual tasks that are largely rules based, structured and repetitive using software robots, also known as bots. RPA tools map a process for a robot to follow which allows the bot to operate in place of a human. Transferring human tasks over to machines. This often involves highly repetitive tasks which robots can do more efficiently. |
| Robotics | A field of AI which incorporates hardware (robots) into the 'thinking' software. |
| Sampling Bias | Is bias in the process of collecting data (e.g., over-sampling or under-sampling of specific groups) |
| Secondary Harm | Refers to any adverse effects experienced by an individual (or organisation) not directly engaged with the AI system, or a subsequent harm identified after an initial harm is experienced by an individual (or organisation) engaged with the AI system. |

| | |
|--|--|
| Significant Harm | Always context specific, a harm which leads to significant concerns. Examples are safety and welfare. |
| Social Intelligence or Affective Computing | Systems that recognize, interpret, or process human affects like emotions, vocal tone, facial expression, etc. |
| Speech Recognition | Relates to the ability of computers to analyse audio files in order to recognise and interpret spoken language. |
| Statistical Bias | Refers mostly to a model that consistently generates an error in prediction when compared with the expected outcome |
| Structured Data | Information with a high degree of organization, such that inclusion in a relational database is seamless and readily searchable by simple, straightforward search engine algorithms or other search operations |
| Supervised Learning | Machine Learning where example inputs to outputs are given for reference to the algorithm. The algorithm learns how to replicate those connections or reach those conclusions. Humans have intervened to direct the algorithm. A form of ML that involves providing the system with data that contains both the input variables (features) and output variables (label, answer) to train the system. |
| Training | In Machine Learning, it is the process of improving the machines performance through running it with more data; usually the algorithm must be given a training dataset that contains the "right" answer for it to learn from initially, but later data sets can omit the answer to evaluate its performance. More and better data will help construct a better model. |
| Turing Test | A test developed by British AI pioneer Alan Turing in the 1950's to determine if a technology is intelligent. Simply it says that if a human interacting with a machine is unable to identify it as a machine and not a human, then that machine is intelligent. This test is often applied to chatbots. |
| Underfitting | Refers to situations in which a Machine Learning system is unable to capture the underlying information contained in the data. |

| | |
|------------------------------|---|
| <p>Unstructured Data</p> | <p>Data that is poorly organized so to be hard for us to understand it or search within it. Data may be unlabelled or uncategorized, or poorly so. A messy email inbox may be an example of unstructured data because while data are labelled by sender and subject, they may be disorganized within the inbox.</p> |
| <p>Unsupervised Learning</p> | <p>Machine learning where no example inputs to outputs or connections are identified for the algorithm. Instead, it independently seeks to identify connections or patterns. An application would be data mining in a large dataset. A form of ML that involves providing the system with data that do not contain the output</p> |
| <p>Weighting</p> | <p>Giving different importance levels to different bits of information. More important information is weighted more heavily. In neural networks, the nodes that perform analysis of different information are weighted based on how much they contribute to the final decisions.</p> |

