



BCS Foundation Certificate in Artificial Intelligence Syllabus

**Version 1.3
November 2023**

This professional certificate is not regulated by the following United Kingdom Regulators -
Ofqual, Qualification in Wales, CCEA or SQA

Contents

Change History	3
Introduction	4
Levels of Knowledge / SFIA Levels	4
Learning Outcomes	4
Target Audience	5
Study Format and Duration	5
Eligibility for Examination	5
Examination Format and Duration	6
Additional Time	6
Guidelines for Accredited Training Organisations	6
Syllabus and Question Weightings	7
Trainer Criteria	7
Classroom Size	7
Invigilator to Candidate Ratio	7
Excerpts from BCS Books	7
Syllabus	8
1. Ethical and Sustainable Human and Artificial Intelligence.....	8
2. Artificial Intelligence and Robotics.....	8
3. Applying the benefits of AI - challenges and risks	9
4. Starting AI how to build a Machine Learning Toolbox - Theory and Practice	9
5. The Management, Roles and Responsibilities of humans and machines	9
Recommended Reading	10
Additional Reading	11
Abbreviations	13
Glossary of Terms	14

Change History

Any changes made to the syllabus shall be clearly documented with a change history log. This shall include the latest version number, date of the amendment and changes made. The purpose is to identify quickly what changes have been made.

Version Number	Changes Made
V1.3 Nov 2023	Updated reading list
V1.2 Aug 2023	SPAG changes throughout, missing publishing information added to reading list.
V1.1 Oct 2020	Amendment to Description. Agent Modelling changed to Intelligent Agent.
V1.0 Oct 2019	Finalised
WIPV8 August 2019	Draft work in progress

Introduction

Artificial Intelligence (AI) is a methodology for using a non-human system to learn from experience and imitate human intelligent behaviour. The BCS Foundation Certificate in Artificial Intelligence tests a candidate's knowledge and understanding of the terminology and general principles of AI. This syllabus covers the potential benefits and challenges of Ethical and Sustainable Robust Artificial Intelligence; the basic process of Machine Learning (ML) – Building a Machine Learning Toolkit; the challenges and risks associated with an AI project, and the future of AI and Humans in work. The Foundation Certificate includes and expands on the knowledge taught in the BCS Essentials Certificate in AI.

Levels of Knowledge / SFIA Levels

This syllabus will provide candidates with the levels of difficulty highlighted within the following table, also enabling them to develop the skills to operate at the highlighted level of responsibility (as defined within the SFIA framework) within their workplace. The levels of knowledge and SFIA levels are further explained on the website www.bcs.org/levels.

Level	Levels of Knowledge	Levels of Skill and Responsibility (SFIA)
7		Set strategy, inspire, and mobilise
6	Evaluate	Initiate and influence
5	Synthesise	Ensure and advise
4	Analyse	Enable
3	Apply	Apply
2	Understand	Assist
1	Remember	Follow

Learning Outcomes

Candidates should be able to demonstrate a knowledge and understanding in the application of Ethical and Sustainable Artificial Intelligence:

1. Human-centric ethical and sustainable human and artificial intelligence.
2. Artificial intelligence and robotics.
3. Applying the benefits of AI projects - challenges and risks.
4. Machine learning theory and practice – building a machine learning toolbox.
5. The management, roles and responsibilities of humans and machines – the future of AI.

Target Audience

The Artificial Intelligence Foundation Certificate is focused on individuals with an interest in, (or need to implement) AI in an organisation, especially those working in areas such as science, engineering, knowledge engineering, finance, education, or IT services.

The following roles could be interested:

- Engineers
- Scientists
- Professional Research Managers
- Chief Technical Officers
- Chief Information Officers
- Organisational Change Practitioners and Managers
- Business Change Practitioners and Managers
- Service Architects and Managers
- Program and Planning Managers
- Service Provider Portfolio Strategists / Leads
- Process Architects and Managers
- Business Strategists and Consultants
- Web Page Developers

Study Format and Duration

Candidates can study for this certificate in two ways:

- Attending an accredited training course. This will require a minimum of 18 hours of study over a minimum of three days.
- Self-study. Self-study resources include online learning and recommended reading (see syllabus Reading List).

Eligibility for the Examination

There are no specific pre-requisites for the entry to the examination, although accredited training is strongly recommended.

Examination Format and Duration

Type	40 Multiple choice questions
Duration	60 Minutes
Pre-requisites	None, but accredited training is highly recommended
Supervised	Yes
Open Book	No (no materials can be taken into the examination room)
Pass Mark	26/40
Calculators	No
Delivery	Digital only

Additional Time

For Candidates Requiring Reasonable Adjustments Due to a Disability.

Please refer to the [reasonable adjustments policy](#) for detailed information on how and when to apply.

For Candidates Whose Language is Not the Language of the Examination

If the examination is taken in a language that is not the candidate's native/official language, then they are entitled to:

- 25% extra time.
- Use their own paper language dictionary (whose purpose is translation between the examination language and another national language) during the examination. Electronic versions of dictionaries will not be allowed into the examination room.

Guidelines for Accredited Training Organisations

Each major subject heading in this syllabus is assigned an allocated percentage of study time. The purpose of this is:

- 1) Guidance on the proportion of content allocated to each section of an accredited course.
- 2) Guidance on the proportion of questions in the exam.

Courses do not have to follow the same order as the syllabus and additional exercises may be included if they add value to the training course.

Syllabus and Question Weighting

Syllabus Area	Knowledge Level	% Syllabus Weighting	Question per paper of 40
Ethical and sustainable human and artificial intelligence.	K1	20%	8
Artificial intelligence and robotics.	K1	20%	8
Applying the benefits of AI - challenges and risks.	K2	15%	6
Starting AI how to build a machine learning toolbox - theory and practice.	K2	30%	12
The management, roles and responsibilities of humans and machines.	K1	15%	6
		100%	40

Trainer Criteria

Criteria	<ul style="list-style-type: none">• Hold the Foundation Certificate in Artificial Intelligence• Have 3 years' experience in related subject (including, but not limited to: data science, high performance computing, scientific computing, data analytics, statistics, mathematics);• Have taught courses professionally.
----------	--

Classroom Size

Trainer to candidate ratio	1:15
----------------------------	------

Invigilator to Candidate Ratio During Examination

Invigilator to candidate ratio	1: 25
--------------------------------	-------

Excerpts from BCS Books

Accredited Training Organisations may include excerpts from BCS books in course materials. To use excerpts from the books, a licence from BCS is required, which will be considered on request by contacting the Head of Publishing at BCS.

Syllabus

Learning Objectives

1. Ethical and Sustainable Human and Artificial Intelligence (20%)

Candidates will be able to:

- 1.1. Recall the general definition of Human and Artificial Intelligence (AI).
 - 1.1.1. Describe the concept of intelligent agents.
 - 1.1.2. Describe a modern approach to Human logical levels of thinking using Robert Dilt's Model.
- 1.2. Describe what are Ethics and Trustworthy AI, in particular:
 - 1.1.1. Recall the general definition of ethics.
 - 1.2.1. Recall that a Human Centric Ethical Purpose respects fundamental rights, principles, and values.
 - 1.2.2. Recall that Ethical Purpose AI is delivered using Trustworthy AI that is technically robust.
 - 1.2.3. Recall that the Human Centric Ethical Purpose Trustworthy AI is continually assessed and monitored.
- 1.3. Describe the three fundamental areas of sustainability and the United Nation's seventeen sustainability goals.
- 1.4. Describe how AI is part of 'Universal Design,' and 'The Fourth Industrial Revolution'.
- 1.5. Understand that ML is a significant contribution to the growth of Artificial Intelligence.
 - 1.5.1. Describe 'learning from experience' and how it relates to Machine Learning (ML) (Tom Mitchell's explicit definition).

2. Artificial Intelligence and Robotics (20%)

- 2.1. Demonstrate understanding of the AI intelligent agent description, and:
 - 2.1.1. list the four rational agent dependencies.
 - 2.1.2. describe agents in terms of performance measure, environment, actuators, and sensors.
 - 2.1.3. describe four types of agent: reflex, model-based reflex, goal-based and utility-based.
 - 2.1.4. identify the relationship of AI agents with Machine Learning (ML).
- 2.2. Describe what a robot is and:
 - 2.2.1. Describe robotic paradigms,
- 2.3. Describe what an intelligent robot is and:
 - 2.3.1. Relate intelligent robotics to intelligent agents.

3. Applying the benefits of AI - challenges and risks (15%)

- 3.1.** Describe how sustainability relates to human-centric ethical AI and how our values will drive our use of AI will change humans, society, and organisations.
- 3.2.** Explain the benefits of Artificial Intelligence by.
 - 3.2.1. list advantages of machine and human and machine systems.
- 3.3.** Describe the challenges of Artificial Intelligence, and give;
 - 3.3.1. general ethical challenges AI raises.
 - 3.3.2. general examples of the limitations of AI systems compared to human systems.
- 3.4.** Demonstrate understanding of the risks of AI project, and:
 - 3.4.1. give at least one a general example of the risks of AI,
 - 3.4.2. describe a typical AI project team in particular,
 - 3.4.2.1. describe a domain expert,
 - 3.4.2.2. describe what is 'fit-of-purpose',
 - 3.4.2.3. describe the difference between waterfall and agile projects.
- 3.5.** List opportunities for AI.
- 3.6.** Identify a typical funding source for AI projects and relate to the NASA Technology Readiness Levels (TRLs).

4. Starting AI how to build a Machine Learning Toolbox - Theory and Practice (30%)

- 4.1.** Describe how we learn from data – functionality, software, and hardware,
 - 4.1.1. List common open-source machine learning functionality, software, and hardware.
 - 4.1.2. Describe introductory theory of Machine Learning.
 - 4.1.3. Describe typical tasks in the preparation of data.
 - 4.1.4. Describe typical types of Machine Learning Algorithms.
 - 4.1.5. Describe the typical methods of visualising data.
- 4.2.** Recall which typical, narrow AI capability is useful in ML and AI agents' functionality.

5. The Management, Roles and Responsibilities of humans and machines (15%)

- 5.1.** Demonstrate an understanding that Artificial Intelligence (in particular, Machine Learning) will drive humans and machines to work together.
- 5.2.** List future directions of humans and machines working together.
- 5.3.** Describe a 'learning from experience' Agile approach to projects
 - 5.3.1. Describe the type of team members needed for an Agile project.

Recommended PRE-COURSE Reading

Title: Human + Machine - Reimagining Work in the Age of AI
Author: Paul R. Daugherty and H. James Wilson,
Publisher: Harvard Business Review Press
Publication Date: 2018
ISBN: 1633693869.

Title: Artificial Intelligence Foundations (learning from experience)
Author: Andrew Lowe and Steve Lawless
Publisher: BCS
Publication Date: 2021
ISBN: 9781780175300

Recommended POST-COURSE Reading

Title: Ethics Guidelines for Trustworthy AI
Author: High-Level Expert Group on Artificial Intelligence
Publisher: European Commission B-1049 Brussels
Publication Date: April 2019.

Title: Artificial Intelligence, A Modern Approach (3rd edition)
Author: Stuart Russell and Peter Norvig
Publisher:
Publication Date: 2016
ISBN 10: 1292153962

Title: Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems
Author: Aurélien Géron
Publisher: O'Reilly
Publication Date: 2017
ISBN 1491962291

Title: The Singularity is Near
Author: Ray Kurzweil
Publisher: Duckworth Overlook
Publication Date: 2005
ISBN: 978-0715635612

Title: The Fourth Industrial Revolution
Author: Klaus Schwab
Publisher: Penguin Random House
Publication Date: 2016
ISBN: 978-0-241-30075-6

Title: Artificial Intelligence Foundations (learning from experience)
Author: Andrew Lowe and Steve Lawless
Publisher: BCS
Publication Date: 2021
ISBN: 9781780175300

Additional Reading – Specialist Reference List

Title: Linear Algebra and Learning from Data (1st edition)
Author: Gilbert Strang
Publisher: Wellesley-Cambridge Press
Publication Date: 2019
ISBN: 13 978-0692196380.

Title: An Introduction to Linear Algebra (5th edition)
Author: Gilbert Strang
Publisher: Wellesley-Cambridge Press
Publication Date: 2016
ISBN: 978-0-9802327-7-6

Title: Novacene: The Coming of Age of Hyperintelligence
Author: James Lovelock
Publisher: Allen Lane - Penguin
Publication Date: 2019
ISBN: 978-0-241-39936-1

Title: The Mystery of Consciousness
Author: John R. Searle
Publisher: The New York Review of Books
Publication Date: 1997
ISBN: 978-0-940322-06-6

The Royal Society: <https://royalsociety.org/topics-policy/projects/machine-learning/>

Title: Machine Learning
Author: Tom Mitchell
Publisher: McGraw-Hill
Publication Date: 1997
ISBN: 0071154671

Title: Life 3.0
Author: Max Tegmark
Publisher: Penguin Books
Publication Date: 2017
ISBN: 978-0-141-98180-2

Title: The Conscious Mind
Author: David Chalmers
Publisher: Oxford University Press
Publication Date: 1996
ISBN: 978-0-19-511789-9

Title: Sustainable Energy – without hot air
Author: Sir David JC Mackay
Publisher: UIT Cambridge Ltd
Publication Date: 2009
ISBN: 978-1-906860-01

Title: Novacene: The Coming of Age of Hyperintelligence
Author: James Lovelock
Publisher: Allen Lane - Penguin
Publication Date: 2019
ISBN: 978-0-241-39936-1

Title: HOW BAD ARE BANANAS? – THE CARBON FOOTPRINT OF EVERYTHING
Author: Mike Berners-Lee
Publisher: Profile Books Ltd
Publication Date: 2010
ISBN: 978-1-84668-891-1

Title: Machine Learning – A Probabilistic Perspective
Author: Kevin P. Murphy
Publisher: MIT
Publication Date: 2012
ISBN: 10:0262018020

Title: Automated Planning Theory and Practice
Author: Malik Ghallab, Dana Nau and Paolo Traverso
Publisher: Elsevier
Publication Date: 2004
ISBN: 1-55860-856-7

Title: The Cambridge Handbook of Artificial Intelligence
Author: Keith Frankish and William Ramsey
Publisher:
Publication Date: 2014
ISBN: 978-0-521-69191-8

Title: Artificial Intelligence: 101 Things You Must Know Today About Our Future
Author: Lasse Rouhiainen
Publisher: CreateSpace Independent Publishing Platform
Publication date: 2018
ISBN: 1982048808.

Title: The Mythical Man Month
Author: Frederick P. Brooks, JR., Addison Wesley
Publisher: Addison-Wesley Professional
Publication Date: 1995
ISBN: 0-201-83595-9

Advanced theoretical Text - <http://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning>

Title: Machine Learning for Absolute Beginners: A Plain English Introduction (2nd edition)
Author: Oliver Theobald
Publisher: Scatterplot Press
Publication Date: 2017
ISBN: 1549617214

Title: Artificial Intelligence Foundations (learning from experience)
Author: Andrew Lowe and Steve Lawless
Publisher: BCS
Publication Date: 2021
ISBN: 9781780175300

Abbreviations

Not all abbreviations need to be used in the course, they are for reference.

Abbreviation	Meaning
AI	Artificial Intelligence
IoT	Internet of Things
ANN	Artificial Neural Network
NN	Neural Network
CNN	Convolution Neural Network
ML	Machine Learning
OCR	Optical Character Recognition
NLP	Natural Language Processing
DL	Deep Learning
DNN	Deep Neural Networks
AGI	Artificial General Intelligent
CPU	Central Processing Unit
GPU	Graphical Processing Unit
RPA	Robotic Process Automation
CART	Classification and Regression Trees
IT	Information Technology
IQ	Intelligence Quotient
EQ	Emotional Quotient

Glossary of Terms

Not all terms need to be used in the course, they are for reference.

Term	Description or Definition
Activation Function	The activation function defines the output of a node given an input or set of inputs.
Intelligent Agent	An intelligent agent (IA) is autonomous, observes through sensors and acts on its environment using actuators.
Algorithm	An algorithm is an unambiguous specification of how to solve a class of problems.
Artificial Intelligence (AI)	A branch of computer science dealing with the simulation of intelligent behaviour in computers.
Automation	Automatically controlled operation of an apparatus, process, or system by mechanical or electronic devices that take the place of human labour.
Autonomous	Undertaken or carried on without outside control
Axon	An axon is a long, slender projection of a nerve cell, or neuron, that typically conducts electrical impulses.
Axon Terminals	Axon terminals are terminations of the <u>telodendria</u> (branches) of an axon.
Back-propagation	A method used in artificial neural networks to calculate a gradient required in the calculation of the weights to be used in the network.
Bayesian Network	A Bayesian network or belief network is a probabilistic graphical model that represents a set of variables and their conditional dependencies.
Bias	Deviation of the expected value of a statistical estimate from the quantity it estimates.
Big Data	Big data is data sets that are so big and complex that traditional data-processing application software are inadequate to deal with them.
Boosting	Boosting is an ensemble meta-algorithm for reducing bias, and also variance in supervised learning and a family algorithms that convert weak learners to strong ones.
Bootstrap Aggregating – Bagging	Bootstrap aggregating, is an ensemble meta-algorithm used in statistical classification and regression.
Chatbot	A chatbot is an artificial intelligence program that conducts a conversation via auditory or textual methods.
Classification	Classification is the problem of identifying to which of a set of classes a new observation belongs.
Clustering	Clustering groups a set of objects in such a way that objects in the same group are more similar to each other than to those in other groups.
Cognitive Simulation	Cognitive simulation uses computers that test how the human mind works.

Term	Description or Definition
Combinatorial Complexity	The exponential growth in computer power required to solve a problem that has many combinations with increasing complexity.
Combinatorial Explosion	A combinatorial explosion is the rapid growth of the complexity of a problem due to the combinations of the problem's input parameters.
Connectionist	Cognitive science that hopes to explain intellectual abilities using artificial neural networks.
Data Analytics	The discovery, interpretation, and communication of meaningful patterns in data.
Data Cleaning	Data cleaning detects and corrects (or removes) corrupt or inaccurate records from a record set, table, or database and refers to identifying incomplete, incorrect, inaccurate or irrelevant parts of the data and then replacing, modifying, or deleting the dirty or coarse data.
Data Mining	The process of discovering patterns in large data sets.
Data Science	Data science uses scientific methods, processes, algorithms, and systems to understand data.
Data Scrubbing	See data cleaning.
Decisions Trees	A decision tree is a decision support tool that uses a tree-like graph or model of decisions and their possible consequences.
Deep Learning	Deep learning is a class of algorithms that use a cascade of multiple layers for feature extraction and transformation. Each successive layer uses the output from the previous layer as input.
Dendrites	Dendrites are branched extensions of a nerve cell that propagate the electrochemical stimulation.
Edges	Edges are the machine learning name for the brain's axons
Ensemble	Ensemble methods use multiple learning algorithms to obtain better predictive performance than could be obtained from any of the constituent learning algorithms alone.
Expert Systems	An expert system is a computer system that emulates the decision-making ability of a human expert.
Feedforward Neural Network	A feedforward neural network is an artificial neural network wherein connections between the nodes do not form a cycle.
Functionality	The tasks that a computer software program is able to do.
Genetic Algorithms	A genetic algorithm (GA) is an algorithm inspired by the process of natural selection.
Hardware	Hardware are the physical parts or components of a computer.
Heuristic	Heuristic is a strategy derived from previous experiences with similar problems.
High Performance Computing – Super Computing	HPC or Supercomputing is a computer with a high level of performance compared to a general-purpose computer
Hyper-parameters	A hyperparameter is a parameter whose value is set before the learning process begins.
Inductive Reasoning	Inductive reasoning makes broad generalisations from specific observations.

Term	Description or Definition
Internet of Things (IoT)	The Internet of Things (IoT) is the network of physical devices, vehicles, home appliances, and other items embedded with electronics, software, sensors, actuators, and connectivity which enables these things to connect and exchange data.
k-Means	<i>k</i> -means is a clustering algorithm that partitions observations into <i>k</i> clusters in which each observation belongs to the cluster with the nearest mean, serving as a prototype of the cluster.
k-Nearest Neighbours	The simplest clustering algorithm used to classify new data points based on the relationship to nearby data points.
Layers	Neural networks are organised into layers and a layer a set of inter-connected nodes.
Linear Algebra	Linear algebra is the branch of mathematics concerning linear equations and functions and their representations through matrices and vector spaces.
Logistic Regression	Logistic Regression is used in binary classification to predict two discrete classes.
Machine Learning (ML)	Machine learning is a subset of artificial intelligence in the field of computer science that gives computers the ability to learn from data.
Model Optimisation	The improvement of the output of a machine learning algorithm (e.g. adjusting hyper parameters)
Natural Language Processing (NLP)	Natural language processing (NLP) is an area of artificial intelligence concerned with the interactions between computers and human (natural) languages, in particular how .
Natural Language Understanding (NLU)	Natural language understanding is term used to describe machine reading comprehension
Nearest Neighbour Algorithm	The nearest neighbour algorithm was one of the first algorithms used to determine a solution to the travelling salesman problem.
Neural Network (NN)	A Machine Learning Algorithm that is based on a mathematical model of the biological brain
Nodes	Nodes represent neurons (biological brain) and are interconnected to form a neural network.
One-hot Encoding	Transforms text-based features into a numerical form, e.g. false is given the number zero and true is given the number 1.
Ontology	Ontology is the philosophical study of the nature of being, becoming, existence, or reality, as well as the basic categories of being and their relations.
Optical Character Recognition (OCR)	Optical character recognition is the conversion of images of typed, handwritten, or printed text into machine-encoded text.
Over-fitting or Over-training	Overfitting is a machine learning model that is too complex, has high variance and low bias. It is the opposite of Underfitting or Under-training.
Probabilistic Inference	Probabilistic Inference uses simple statistical data to build nets for simulation and models.
Probability	Probability is the measure of the likelihood that an event will occur.
Pruning	Pruning reduces the size of decision trees.

Term	Description or Definition
Python	A programming language popular in machine learning
Random Decision Forests	Random decision forests are an ensemble learning method for classification, regression and other tasks.
Random Forests	Random forests are an ensemble learning method or classification, regression and other tasks, that operate by constructing a multitude of decision trees at training time.
Regression Analysis	In machine learning, regression analysis is a simple supervised learning technique used to find a trendline to describe the data.
Reinforcement Machine Learning	Reinforcement learning (RL) uses software agents that take actions in an environment so as to maximize some notion of cumulative reward.
Robotics	Robotics deals with the design, construction, operation, and use of robots, as well as computer systems for their control, sensory feedback, and information processing.
Robotic Process Automation (RPA)	Robotic process automation is a business process automation technology based on the notion of <i>software robots</i> or artificial intelligence workers.
Scripting	Scripting are programs written for a special run-time environment that automate the execution of tasks that could alternatively be executed one-by-one by a human operator.
Search	The use of machine learning in search problems, e.g. shortest path
Semi-supervised Machine Learning	Machine learning that uses labelled and unlabelled data for training.
Sigmoid Equation	A sigmoid function is a mathematical function having a characteristic "S"-shaped curve or sigmoid curve.
Software	Software is a generic term that refers to a collection of data and computer instructions that tell the computer how to work.
Software Robots	A software robot replaces a function that a human would otherwise do.
Strong AI or Artificial General Intelligence	Strong AI's goal is the development of artificial intelligence to the point where the machine's intellectual capability is functionally equal to a human.
Supervised Machine Learning	Supervised machine learning is the task of learning a function that maps an input to an output based on example input-output pairs.
Support Vector Machine	A support vector machine constructs a hyperplane or set of hyperplanes in a high- or infinite-dimensional space, which can be used for classification, regression, or other tasks like outliers detection.
Swarm-intelligence	Swarm intelligence is the collective behaviour of decentralized, self-organised systems, natural or artificial
Symbolic	Symbolic artificial intelligence is the term for the collection of all methods in artificial intelligence research that are based on high-level "symbolic" (human-readable) representations of problems, logic and search.
System	A regularly interacting or interdependent group of items forming a unified whole.

Term	Description or Definition
The Fourth Industrial Revolution	The Fourth Industrial Revolution builds on the Digital Revolution, representing new ways in which technology becomes embedded within societies and even the human body.
Turing Machine	A Turing machine is a mathematical model of computation.
Un-supervised Machine Learning	Unsupervised machine learning infers a function that describes the structure of "unlabelled" data.
Under-fitting	Underfitting is when the machine learning model has low variance and high bias. It is the opposite of Over-fitting or Over-training.
Universal Design	Universal design (close relation to inclusive design) refers to broad-spectrum ideas meant to produce buildings, products and environments that are inherently accessible to older people, people without disabilities, and people with disabilities.
Validation Data	A set of data used to test the output of a machine learning model that is not used to train the model.
Variance	Variance is the expectation of the squared deviation of a random variable from its mean.
Visualisation	Visualisation is any technique for creating images, diagrams, or animations to communicate a message.
Weak AI or Narrow AI	Weak artificial intelligence (weak AI), also known as narrow AI, is artificial intelligence that is focused on one narrow task. It is the contrast of Strong AI.
Weights	A weight function is a mathematical device used when performing a sum, integral, or average to give some elements more "weight" or influence on the result than other elements in the same set.