



BCS Essentials Certificate in Artificial Intelligence Syllabus V1.0

Change History

This log provides a single point of reference, where a summary of any changes is recorded, to include the date of the amendment and a summary of the changes made.

Version Number	Changes Made
Version 1.0 February 2019	Signed off; made live
Version 0.6 January 2019	Reading list updated
Version 0.5 December 2018	Added additional LOs: 2.4, 2.5, 3.3, 3.4, 3.5
Version 0.4 October 2018	BCS Formatted syllabus created.

Contents

Change History	2
Introduction	4
Levels of Knowledge / SFIA Levels	4
Learning Outcomes.....	4
Target Audience.....	5
Course Format and Duration.....	5
Examination Format and Duration.....	6
Eligibility for the Examination	6
Additional Time	6
Guidelines for Accredited Training Organisations.....	7
Syllabus	8
Learning Objectives	8
1 Artificial and Human Intelligence: an Introduction and History (25%).....	8
2 Examples of AI: Benefits, Challenges and Risks (30%)	8
3 An introduction to Machine Learning (35%).....	8
4 The Future of Artificial Intelligence – Human and Machine Together (10%)	9
Exam Weighting.....	9
Recommended Reading List	10
Abbreviations	12
Glossary of Terms.....	12

Introduction

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

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](#).

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 basic knowledge and understanding of general concepts in the following areas:

1. Human and Artificial Intelligence;
2. The Machine Learning process;
3. The benefits, challenges and risks of a Machine Learning project;
4. The future of humans and machines in Work.

Target Audience

The Artificial Intelligence Essentials certificate is focussed 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, 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

Course Format and Duration

Candidates can choose to study for this certificate from one of two ways: by either attending a training course provided by a BCS Accredited Training Organisation, or by self-study.

BCS recommends that for full coverage of the syllabus to be achieved, training courses leading to the certificate should normally run for a minimum 6 hours over 1 day.

Candidates should spend about 40 hours on self-study with no prior knowledge.

Examination Format and Duration

Type	20 Multiple choice questions
Duration	30 Minutes
Supervised	Yes
Open Book	No
Pass Mark	13/20
Calculators	No calculators nor mobile technology is permitted
Delivery	Paper based examination

Eligibility for the Examination

There are no specific pre-requisites for the entry to the examination, although a good knowledge of computers and a basic understanding of data using computers (e.g. spreadsheets) is highly recommended.

Additional Time

For Candidates Requiring Reasonable Adjustments Due to a Disability

Please refer to the [reasonable adjustments policy](#) for 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 time 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.

Trainer Criteria

The following criterion apply:

- Hold a Foundation Certificate;
- 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

Trainers may instruct up to 15 candidates.

Invigilator to Candidate Ratio During Examination

Up to 25 candidates per invigilator.

Syllabus

Learning Objectives

1. Artificial and Human Intelligence: An Introduction and History (25%)

Candidates will be able to:

- 1.1 Recall the general definition of human and Artificial Intelligence (AI);
- 1.2 Describe 'learning from experience' and how it relates to Machine Learning (ML) (Tom Mitchell's explicit definition);
- 1.3 Understand that ML is a significant contribution to the growth of Artificial Intelligence;
- 1.4 Describe how AI is part of '*Universal Design*,' and '*The Fourth Industrial Revolution*'.

2. Examples of AI: Benefits, Challenges and Risks (30%)

Candidates will be able to:

- 2.1 Explain the benefits of Artificial Intelligence, and
 - 2.1.1 list advantages of machine and human and machine systems;
- 2.2 Describe the challenges of Artificial Intelligence, and give:
 - 2.2.1 general examples of the limitations of AI compared to human systems,
 - 2.2.2 general ethical challenges AI raises.
- 2.3 Demonstrate understanding of the risks of Artificial Intelligence, and
 - 2.3.1 give at least one a general example of the risks of AI;
- 2.4 Identify a typical funding source for AI projects;
- 2.5 List opportunities for AI.

3. An introduction to Machine Learning (35%)

Candidates will be able to:

- 3.1 Demonstrate understanding of the AI intelligent agent description, and:
 - 3.1.1 identify the differences with Machine Learning (ML), and:
 - 3.1.2 list the four rational agent dependencies,
 - 3.1.3 describe agents in terms of performance measure, environment, actuators and sensors,

- 3.1.4 describe four types of agent: reflex, model-based reflex, goal-based and utility-based.
- 3.2 Give typical examples of Machine Learning in the following contexts:
 - 3.2.1 business,
 - 3.2.2 social (media, entertainment),
 - 3.3.3 science.
- 3.3 Recall which typical, narrow AI capability is useful in ML and AI agents' functionality;
- 3.4 Describe and give examples of the following forms of ML:
 - 3.4.1 supervised,
 - 3.4.2 unsupervised,
 - 3.4.3 reinforcement.
- 3.5 Describe the basic schematic of a neural network.

4. The Future of Artificial Intelligence – Human and Machine Together (10%)

Candidates will be able to:

- 4.1 Demonstrate an understanding that Artificial Intelligence (in particular, Machine Learning) will drive humans and machines to work together;
- 4.2 List future directions of humans and machines working together.

Exam Weighting

Syllabus Area	Target Number of Questions
Artificial and Human Intelligence: An Introduction and History	25% - 5 Questions – K1
Examples of AI - Benefits, Challenges and Risks	30% - 6 Questions – K2
An introduction to Machine Learning	35% - 7 Questions – K1
The Future of Artificial Intelligence – Human and Machine Together	10% - 2 Questions – K1

Recommended Reading List

Artificial Intelligence and Consciousness

Title Artificial Intelligence, A Modern Approach, 3rd Edition
Author Stuart Russell and Peter Norvig,
Publication Date 2016,
ISBN 10 1292153962

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

Title The Conscious Mind
Author David Chalmers
Publication Date 1996
ISBN 978-0-19-511789-9

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

Machine Learning

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

Title Machine Learning For Absolute Beginners: A Plain English Introduction (2nd Edition),
Author Oliver Theobald,
Publication Date 2017
ISBN 1549617214.

High Level / Management Consultant View of Machine Learning

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

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.

High Level / Research and Political View of Machine Learning

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

Professional Development of Machine Learning Algorithms and Planning

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

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

Additional Reading:

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

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

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.

Abbreviations

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

Term	Description or Definition
Activation Function	The activation function defines the output of a node given an input or set of inputs.
Agent Modelling	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 telodendria (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.

Term	Description or Definition
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 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.
Combinatorial Complexity	Is 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
Emotional Intelligence or Emotional Quotient (EQ)	The understanding of our emotions and the emotions of others.
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 can do
Genetic Algorithms	a genetic algorithm (GA) is an algorithm inspired by the process of natural selection.

Term	Description or Definition
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.
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.
Intelligent Quotient (IQ)	Is a standard test of intelligence.
k-Means	k-Means is a clustering algorithm that partitions observations into k clusters, where 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 is a set of interconnected 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, and how these happen.
Natural Language Understanding (NLU)	Natural Language Understanding is the 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 Under-fitting or Under-training.
Probabilistic Inference	Probabilistic Inference uses simple statistical data to build nets for simulation and models.

Term	Description or Definition
Probability	Probability is the measure of the likelihood that an event will occur.
Pruning	Pruning reduces the size of decision trees.
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 for 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 (RL)	Reinforcement Learning (RL) uses software agents that take actions in an environment to maximise 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 Function	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's.
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 outlier detection.
Swarm Intelligence (SI)	Swarm Intelligence is the collective behaviour of decentralised, 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.
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.

Term	Description or Definition
Unsupervised Machine Learning	Unsupervised Machine Learning infers a function that describes the structure of unlabelled data.
Underfitting	Underfitting is when the machine learning model has low variance and high bias. It is the opposite of Overfitting or Overtraining.
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.