Short Course in Artificial Intelligence

Sample Materials

1. Course Overview and Objectives

This Artificial Intelligence (AI) short course will provide learners with the foundation and branches of Artificial Intelligence, such as search strategies, knowledge representation and reasoning techniques. This course will also introduce a range of well-known techniques and applications in artificial intelligence, such as fuzzy logic machine learning, expert systems, natural language processing, and intelligent agents.

At the end of the course, learners should be able to:

1. Understand the importance of AI and its applications.
2. Apply a range of well-established AI search strategies and knowledge representation techniques in problem solving.
3. Assess a range of well-established techniques for reasoning with uncertain knowledge.
4. Understand a range of machine learning techniques.
5. Understand the range of AI techniques that are being applied in industry or research.
6. Implement and evaluate a range of AI models and techniques for solving real-world problems.

2. Learning Outcomes and Objectives

Learning Outcomes; The Learner will:	**Objectives;** The Learner can:
1. Understand the importance of AI and its applications.	1.1 Explain the meaning of AI and its origin. 1.2 Identify the characteristics of AI. 1.3 Discuss the limitations and ethics of AI. 1.4 Discuss current and future developments in the field of AI and its applications.
2. Apply a range of well-established AI search strategies and knowledge representation techniques in problem solving.	2.1 Construct simple state spaces. 2.2 Select and apply appropriate search techniques in problem solving. 2.3 Explain and identify different types of knowledge. 2.4 Apply knowledge representation using the logical, semantic network, frame, and production rules techniques.
3. Assess a range of well-established techniques for reasoning with uncertain knowledge.	3.1 Explain the concept of uncertainty. 3.2 Explain the source of uncertain knowledge. 3.3 Discuss and apply probabilistic reasoning using the Bayes' rule and the certainty factor theory.
4. Understand a range of machine learning techniques.	4.1 Explain machine learning. 4.2 Compare and contrast supervised learning, unsupervised learning, and reinforcement learning. 4.3 Identify and apply suitable machine learning techniques in problem solving.

5. Understand the range of AI techniques that are being applied in industry or research.	5.1 Describe the key players, components, characteristics, and limitations of expert systems. 5.2 Apply reasoning techniques in rule-based expert systems. 5.3 Explain natural language processing and its components. 5.4 Discuss a range of applications using natural language processing. 5.5 Explain intelligent agents and PEAS. 5.6 Discuss a range of applications using intelligent agents.
6. Implement and evaluate a range of AI models and techniques for solving real-world problems.	6.1 Select suitable tools and techniques for use in designing AI models. 6.2 Construct an AI model. 6.3 Evaluate the performance of AI models

3. Syllabus

Syllabus			
Topic No	Title	Proportion	Content
1	Introduction to AI	1/12 2.5 hours of lectures 2.5 hours of exercises	• Definitions • History of AI • Characteristics of AI • Limitations of AI • Ethics in AI • Current and future development of AI and its applications ***Learning Outcome: 1***
2	Problem Solving Using Search	1/12 2.5 hours of lectures 2.5 hours of exercises	• Problem representation in state space • Strategies for state space search • Uninformed search (blind search) • Informed search (heuristic Search) ***Learning Outcome: 2***
3	Knowledge Representation	1/12 2.5 hours of lectures 2.5 hours of exercises	• Types of knowledge • Logical representation • Semantic network representation • Frame representation • Production Rules ***Learning Outcome: 2***
4	Uncertain Knowledge	1/12 2.5 hours of lectures 2.5 hours of exercises	• Uncertainty and its sources • Basic probability • Bayes' rule • Bayesian reasoning • Certainty factors ***Learning Outcome: 3***
5	Fuzzy Logic	1/12 3 hours of lectures 3 hours of exercises	• Fuzzy logic • Linguistic variables • Fuzzy sets and operations • Fuzzy rules • Fuzzy system ***Learning Outcome: 3***
6	Machine Learning	1/12 3 hours of lectures 3 hours of exercises	• Introduction • Supervised learning • Unsupervised learning • Reinforcement learning • Applications of machine learning ***Learning Outcome: 4***

7	Neural Networks	1/12 2 hours of lectures 2 hours of exercise 4 hours of laboratory sessions	• Basic structure • Perceptron • Multilayer neural networks • Backpropagation learning • Accelerated learning • Recurrent neural networks • ***Learning Outcome: 4, 5***
8	Decision Trees	1/12 3 hours of lectures 2 hours of exercise sessions 5 hours of laboratory sessions	• Structure and terminologies • Attribute selection measures • ***Learning Outcome: 4, 6***
9	Genetic Algorithms	1/12 3 hours of lectures 2 hours of exercise sessions	• Simulation of Natural Evolution • Basic genetic algorithms • ***Learning Outcome: 4***
10	Expert Systems	1/12 3 hours of lectures 2 hours of exercise sessions 5 hours of laboratory sessions	• The development team of an expert system • Components of an expert system • Characteristics of an expert system • Rule-based expert system • ***Learning Outcome: 5, 6***
11	Natural Language Processing	1/12 3 hours of lectures 2 hours of exercise sessions 5 hours of laboratory sessions	• Terminologies • Components of natural language processing • Phases in natural language processing • Natural language processing pipeline • Applications of natural language processing • ***Learning Outcome: 5, 6***

12	Intelligent Agents	1/12 3 hours of lectures 2 hours of exercise sessions	• Agents and environments • Rationality • PEAS • Types of intelligent agents • Game Playing • Algorithms in games • ***Learning Outcome: 5***

4. Training Delivery Schedule

NCC Education Short Course in Artificial Intelligence contains 12 topics, 8 topics are theoretical topics and 4 topics are practical topics which require additional software to be installed.

As a training provider it is entirely at your discretion how you wish to deliver the content of this course. Training can be delivered following the sequence of the training manual from Topic 1 to Topic 12. In response to the current pandemic situation, training providers may wish to consider opting for a blended training approach by initially covering the 8 theoretical topics remotely first, followed by the 4 practical topics in a face-to-face environment.

Theoretical Topics

Topic No	Title
1	Introduction to AI
2	Problem Solving Using Search
3	Knowledge Representation
4	Uncertain Knowledge
5	Fuzzy Logic
6	Machine Learning
9	Genetic Algorithms
12	Intelligent Agents

Practical Topics

Topic No	Title
7	Neural Networks
8	Decision Trees
10	Export System
11	Natural Language Processing

5. Related National Occupational Standards

The UK National Occupational Standards describe the skills that professionals are expected to demonstrate in their jobs in order to carry them out effectively. They are developed by employers and this information can be helpful in explaining the practical skills that learners have covered in this course.

Related National Occupational Standards (NOS)
Sector Subject Area: IT and Telecoms **Related NOS:** TECIS806401 Develop and communicate data visualisations TECIS806301 Assist in developing data visualisations TECIS804301 Assist in the deployment of artificial intelligence solutions TECIS804401 Manage the deployment of artificial intelligence solutions TECIS805401 Develop and implement machine learning algorithms TECIS805301 Assist in developing and validating machine learning solutions ESKITP4022/3/4/5/6 Data Analysis Levels 2/3/4/5/6 Role

6. Resources

Trainer Guide: This guide contains notes for trainers on the organisation of each topic, and suggested use of the resources. It also contains all of the suggested exercises and model answers.

PowerPoint Slides: These are presented for each topic for use in the training. They contain many examples which can be used to explain the key concepts. Handout versions of the slides are also available.

Learner Guide: This contains the topic overviews, slides screenshots and all of the suggested exercises. The orientation and sequence of the slide's screenshots are shown as below:

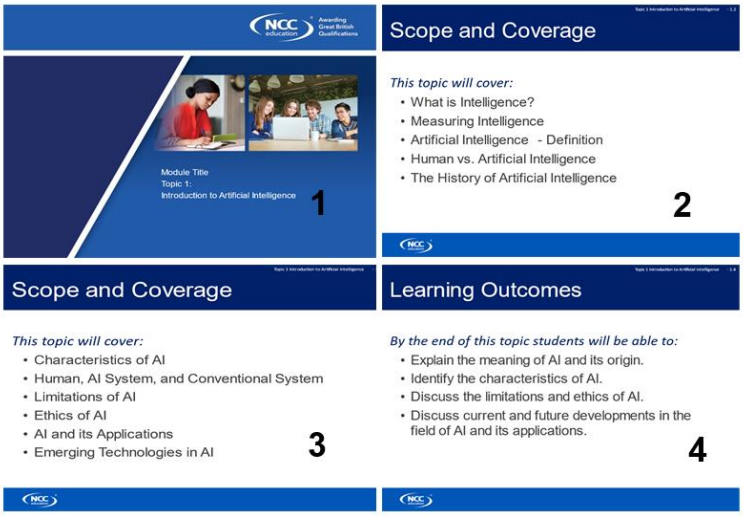

Each learner will need access to this and should bring it to all of the taught hours for the module.

6.1 Additional Hardware and Software Requirements

Hardware: A computer with 64-bit Windows/Mac OS/ Linux (Desktop/Laptop)

Software: Chapter 7: Waikato Environment for Knowledge Analysis (WEKA)
https://waikato.github.io/weka-wiki/downloading_weka/

Chapter 8: Anaconda Python (included Jupyter Notebook, Scikit-learn)
https://www.anaconda.com/products/individual

Chapter 10: SWI-Prolog https://www.swi-prolog.org/download/stable

Chapter 11: Anaconda Python (included Jupyter Notebook) pip install NLTK
https://www.anaconda.com/products/individual

7. Pedagogic Approach

Guided Learning Hours			Total
Lecture	**Exercise**	**Laboratory**	
33	28	19	80

The trainer-led time for this course is comprised of lectures, laboratory sessions and exercises. The breakdown of the hours is also given at the start of each topic, with a minimum of 5 hours of contact time per topic.

7.1 Lectures

Lectures are designed to introduce learners to each topic; PowerPoint slides are presented for use during these sessions. Learners should also be encouraged to be active during this time and to discuss and/or practice the concepts covered. Trainers should encourage active participation and field questions wherever possible.

7.2 Exercises

Exercises provide tasks to involve group work, investigation and independent learning for certain topics. The details of these tasks are provided in this guide and also in the Learner Guide. They are also designed to deal with the questions arising from the trainers and laboratory sessions.

7.3 Laboratory Sessions

During these sessions, learners are required to work through practical exercises. The details of these are provided in this guide and also in the Learner Guide. Some sessions will require more support than others as well as IT resources. More detail is given in this guide.

8. Further Reading List

A selection of sources of further reading around the content of this module must be available in your Accredited Partner Centre's library. The following list provides suggestions of some suitable sources:

- Russell, S., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach*: Pearson
- Negnevitsky, M. (2011). *Artificial Intelligence: A Guide to Intelligent Systems*: Pearson Education Limited.
- Luger, G. F. (2011). *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*: Pearson Education.
- Poole, D.L. and Mackworth, A.K. (2017). *Artificial Intelligence: foundations of computational agents*: Cambridge University Press.
- Giarratano, J. C., & Riley, G. (2005). *Expert Systems: Principles and Programming*: Thomson Course Technology
- Sharda, R., Delen, D., & Turban, E. (2019). *Analytics, Data Science, & Artificial Intelligence: Systems for Decision Support*: Pearson.

Topic 1: Introduction to Artificial Intelligence

1.1 Learning Objectives

This topic provides an overview of Artificial Intelligence (AI). It helps learners to understand the meaning of intelligence in general, way to measure intelligence, the meaning of AI, and the comparison between human intelligence and machine intelligence. A brief history of AI highlighting several key ideas, events, and researches will also be covered to develop appreciation and understanding in the motivation and background of AI.

To further discover what AI can do, the characteristics, limitations, and ethics of AI are discussed. This topic also presents an overview of some of the important applications and emerging technologies in AI.

On completion of the topic, learners will be able to:

- Explain the meaning of AI and its origin.
- Identify the characteristics of AI.
- Discuss the limitations and ethics of AI.
- Discuss current and future developments in the field of AI and its applications.

1.2 Pedagogic Approach

Information and theory of the topic will be presented to the learners during lectures. They will then practise the skills during the exercise sessions.

1.3 Timings

Lectures:	2.5 hours
Exercises:	2.5 hours

1.4 Trainers Notes

The following is an outline of the material to be covered during the lecture time and should be read in conjunction with the slides provided.

The structure of this topic is as follows:

- What is Intelligence?
- Measuring Intelligence
- Artificial Intelligence - Definition
- Human vs. Artificial Intelligence
- The History of Artificial Intelligence
- Characteristics of AI
- Human, AI System, and Conventional System
- Limitations of AI
- Ethics of AI
- AI and its Applications
- Emerging Technologies in AI

1.4.1 Guidance on the Use of the Slides

Slide 5: This slide gives the definitions of intelligence taken from dictionaries. Highlight on:

- the abilities associated with intelligence (e.g. learning ability, understanding ability, problem solving ability, making decision ability). Ask learners to reflect these abilities in human or themselves. Do they see these abilities in other beings or things?
- Philosophers have been trying for over 2000 years to understand and resolve two Big Questions of the Universe: E.g. **How does a human mind work, and Can non-humans have minds?** These questions are still unanswered.
- In a nutshell, the goal of **AI** as a science is to make machines do things that would require intelligence if done by humans.

Slides 6-7: These slides introduce an importance figure in AI, Alan Turing. Introduce Alan Turing and highlight the significant of his paper ***"Computing Machinery and Intelligence"***, which brought to the questions of ***Is there thought without experience? Is there mind without communication? Is there language without living? Is there intelligence without life?*** All these questions, as you can see, are just variations on the fundamental question of artificial intelligence, ***Can machines think?***

Next, the Turing Imitation Game (or known as Turing Test) as a way to measure intelligence in machine. Turing did not provide definitions of machines and thinking, he just avoided semantic arguments by inventing a game. The details of the game can be found in the references.

Slides 8-10: These slide gives definitions of AI by some of the key founders of AI (John McCarthy, Marvin Minsky). Ask learners what they can understand from the definitions. There is also an illustration of common intelligence found in Human and similar or corresponding technology in AI.

Slides 11-25: We now move on to a brief history of AI. The chronology given is according to different phases and era of AI from its emergence, birth, rise, unfulfilled promises, technology of expert systems, neural networks, and up to recent. The main objective is for learners to know the origin of AI and the development of several techniques by some of the key people. The details of the history can be found in references and other online resources. learners can be asked to give other examples of event/idea/creation/development of AI in the past/recent that they found inspiring or amazed.

Slide 26: Checkpoint summary

Slides 27-29: The first slide gives some of the common characteristics found in an AI system/software/solution. Learners will have to understand and familiarise with these terms. The next slide gives a comparison between human, AI system and conventional system (regular system). The objective is for learners to be able to differentiate AI system from conventional system. They should be aware of what are the behaviours of an AI system. Some of the terms such as heuristic, reasoning, inexact and uncertain data may need further explanation.

Awarding Great British Qualifications

Module Title
Topic 1:
Introduction to Artificial Intelligence

Scope and Coverage

This topic will cover:

- What is Intelligence?
- Measuring Intelligence
- Artificial Intelligence - Definition
- Human vs. Artificial Intelligence
- The History of Artificial Intelligence

Scope and Coverage

This topic will cover:

- Characteristics of AI
- Human, AI System, and Conventional System
- Limitations of AI
- Ethics of AI
- AI and its Applications
- Emerging Technologies in AI

Learning Outcomes

By the end of this topic students will be able to:

- Explain the meaning of AI and its origin.
- Identify the characteristics of AI.
- Discuss the limitations and ethics of AI.
- Discuss current and future developments in the field of AI and its applications.

What is Intelligence?

The ability to learn, understand, and make judgments or have opinions that are based on reason
- Cambridge Dictionary

The ability to acquire and apply knowledge and skills.
- Oxford Dictionary

Measuring Intelligence

" I propose to consider the question, Can Machines Think? "
- Alan Turing

- Turing did not provide definitions of machines and thinking, he just avoided semantic arguments by inventing a game, the **Turing Imitation Game**.

Measuring Intelligence with Turing Imitation Game

Phase 1

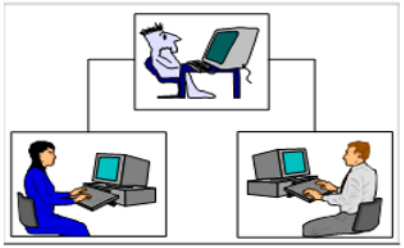

Phase 2

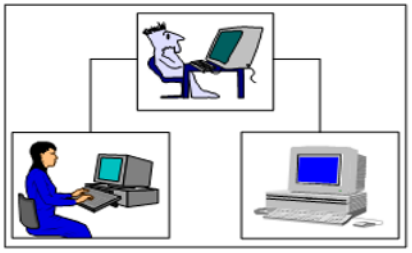

(Negnevitsky, 2011)

Turing believed that by the end of the 20th century it would be possible to program a digital computer to play the imitation game. Although modern computers still cannot pass the Turing test, it provides a basis for the verification and validation of knowledge-based systems.

Artificial Intelligence - Definition

"The science and engineering of making intelligent machines, especially intelligent computer programs."
- John McCarthy, founder of AI

"The science of making machines do things that would require intelligence if done by men."
- Marvin Minsky, founder of AI

1.5 Exercise Sessions

The time allocation for this topic is 2.5 hours.

Trainers' Notes:

Learners have copies of the exercises in the Learner Guide. Answers are not provided in their guide.

Exercise 1:

Define in your own words:

(a) intelligence,
(b) artificial intelligence,
(c) reasoning,
(d) knowledge,
(e) ethics.

Suggested Answer:

(a) intelligence
Dictionary definitions of intelligence are all reasonable answers, but if we want something quantifiable we would use something like - *the ability to learn and understand, to solve problems and to make decisions*.

(b) artificial intelligence
A science to make machines do things that would require intelligence if done by humans.

(c) reasoning
The ability to solve problems through logical deduction. In artificial intelligence, reasoning is essential so that the machine can think rationally as a human to find valid conclusions.

(d) knowledge
The information about a domain that can be used to solve problems in that domain.

(e) ethics
Ethics of AI can be divided into the concern with the moral behaviour of humans as they design, construct, use and treat artificially intelligent agents, and the concern with the moral behaviour of artificially intelligent agents, such as the issues of singularity and superintelligence.

7.6 Laboratory Sessions

The laboratory time allocation for this topic is 4 hours.

Trainers' Notes:

Learners have copies of the laboratory exercises in the Learner Guide. Answers are not provided in their guide. You will need to provide support and advice for learners while they undertake Exercises 1 and 2.

In the exercises learners will explore on Weka, which is open source machine learning software that can be accessed through a graphical user interface, standard terminal applications, or a Java API. It is widely used for teaching, research, and industrial applications, contains a plethora of built-in tools for standard machine learning tasks

Exercise 1:

- Download and install WEKA: https://waikato.github.io/weka-wiki/downloading_weka/
- Download the **manual** for the installed WEKA version and learn up the basic features and WEKA's experimenter application (required for Exercise 2). https://waikato.github.io/weka-wiki/documentation/#general-documentation

Exercise 2:

For this exercise, you will use **WEKA** to learn about the **neural network classifier** (called **MultilayerPerceptron**) and perform a statistical comparison of this classifier to a decision tree classifier (called **J48**) by using several datasets. Answer all the questions below:

1. Run the MultilayerPerceptron classifier on the *weather.arff* dataset. Use the default parameter settings and use the training set as the test option. Include in your submission the printed results (weights and statistics) from WEKA. Show the graphic of the neural network topology (input nodes, hidden nodes, output nodes, connections) used by the classifier.

2. WEKA's default parameter settings (among others) for MultilayerPerceptron are -L 0.3 -M 0.2 -N 500 -H a. Explain in your own words what these means.

Suggested Answers:

1.

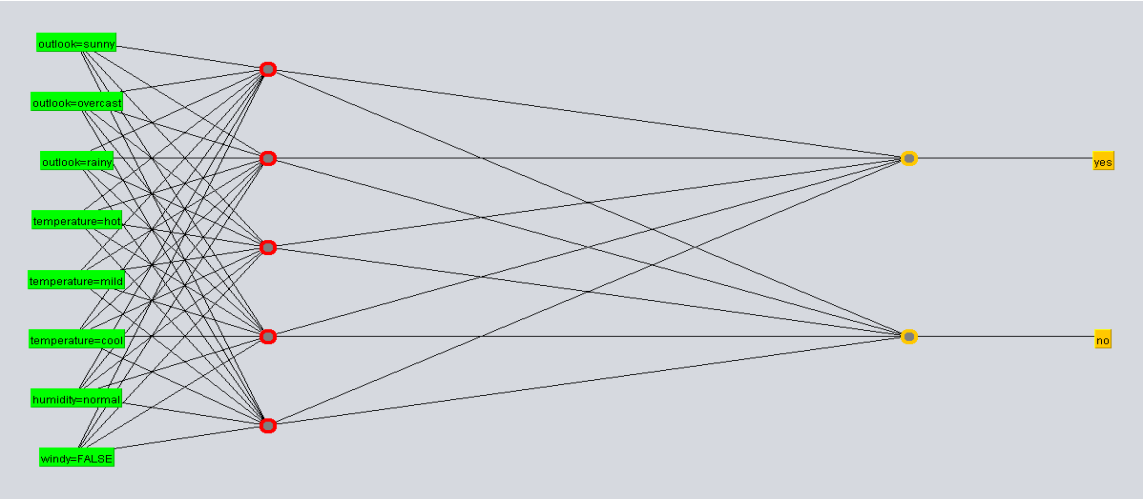

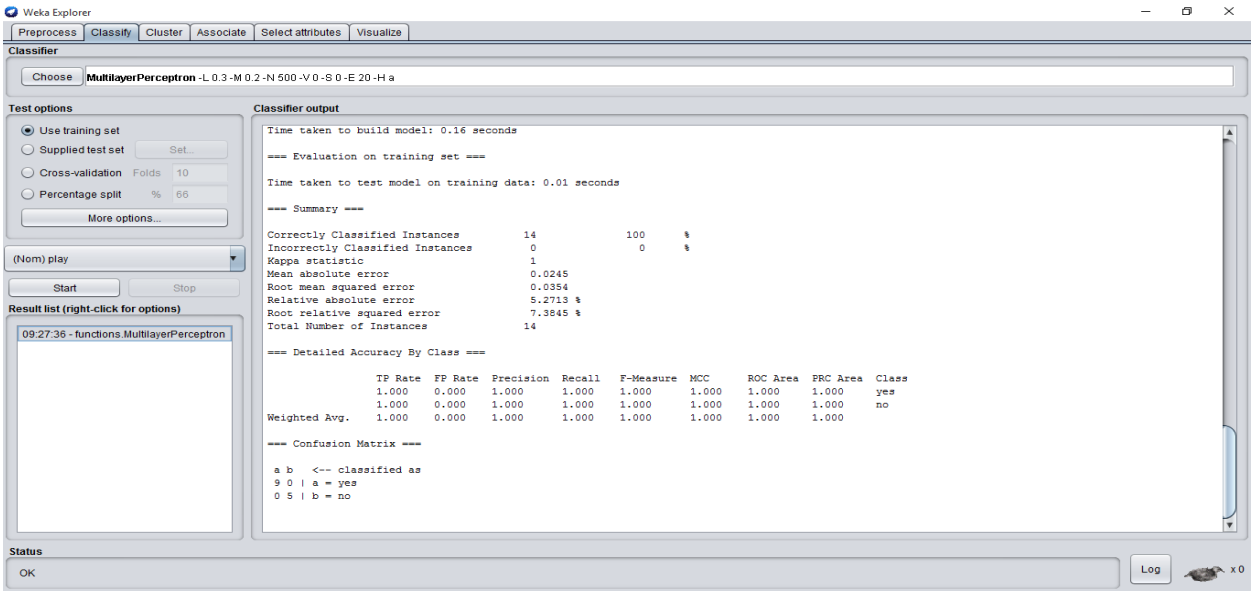

2. MultilayerPerceptron -L 0.3 -M 0.2 -N 500 -H a can be defined as 0.3 for learning rate, 0.2 for momentum, 500 epoch that are complete set of data and hidden layer is a.

Exercise 2: Basic Prolog Syntax

Before we start writing Prolog code you need to understand some of its basic syntax. Like other programming languages it can get complex but we are not going to go in that depth, but this exercise will give you the right confidence to create simple expert systems. Read the following explanations on Prolog.

(a) A Prolog program

- A PROLOG program is a set of **facts** and **rules**.
- A **fact** is basically a thing that is known or proved to be true. For example, *Alice is the parent of Jim*, can be written in prolog as:

```
parent(alice, jim).
```

Notes: Each statement in Prolog should end with dot (.). Anything that starts with a lower-case letter in the parentheses is a literal constant (a specific object or value). Relation name (e.g. parent) should start with a lower case too.

- A simple program may contain just facts as follows:

```
parent(alice, jim).
parent(jim, tim).
parent(jim, dave).
parent(jim, sharon).
parent(tim, james).
parent(tim, thomas).
female(alice).
```

- A **rule** is a thing which are true provided that something else is true. Rules are made from facts. For example, *X is the mother of Y, if X is the parent of Y and X is a female*:

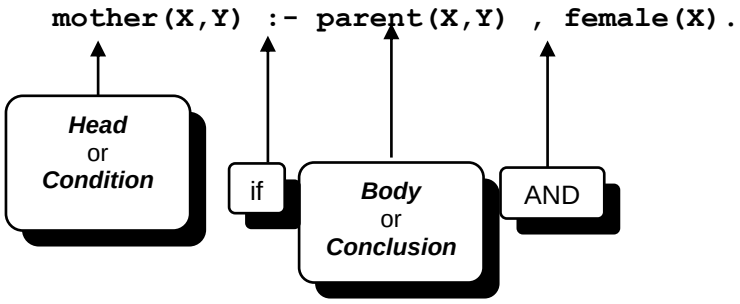

*Notes: Any word written with a capital letter is a **variable.** So **variables** are used to define generic rules. Semicolon ; represents OR.*

End of Document