

2007 Paper 8 Question 9

Artificial Intelligence II

An agent can exist in a state $s \in S$ and can move between states by performing actions, the outcome of which might be uncertain.

- (a) Explain what is meant by a *Utility Function* within this context. [2 marks]
- (b) Give a definition of *Maximum Expected Utility* and describe the way in which it can be used to decide which action to perform next. [3 marks]
- (c) What difficulties might you expect to have to overcome in practice in order to implement such a scheme? [3 marks]
- (d) Explain why it makes sense to use a utility function in the design of an agent, even though it can be argued that real agents (such as humans) appear not to do this, but rather to act on the basis of *preferences*. [4 marks]
- (e) As well as actions allowing an agent to move between states, an agent might be capable of performing actions that allow it to discover more about its environment. Give a full derivation of the *Value of Perfect Information*, and explain how this idea can be used as the basis for an agent that can gather further information in a way that takes account of the potential cost of performing such actions. [8 marks]