

Artificial Intelligence: foundations of computational
agents, Second Edition
Solution Manual

Do not redistribute or create a public link to this
document.

David Poole and Alan Mackworth

June 14, 2017

Contents

Contents	2
1 Artificial Intelligence and Agents	5
2 Agent Architectures and Hierarchical Control	7
3 Searching for Solutions	11
4 Reasoning With Constraints	17
5 Propositions and Inference	31
6 Planning With Certainty	47
7 Supervised Machine Learning	55
8 Reasoning With Uncertainty	65
9 Planning With Uncertainty	81
10 Learning With Uncertainty	97
11 Multiagent Systems	99
12 Learning to Act	103
13 Individuals and Relations	111

<i>Contents</i>	3
14 Ontologies and Knowledge-based systems	121
15 Relational Planning, Learning, and Probabilistic Reasoning	129
16 Retrospect and Prospect	137
Bibliography	139

Chapter 1

Artificial Intelligence and Agents

Exercise 1.1 For each of the following, give five reasons why:

- (a) A dog is more intelligent than a worm.
- (b) A human is more intelligent than a dog.
- (c) An organization is more intelligent than an individual human.

Based on these, give a definition of what “more intelligent” may mean.

Exercise 1.2 Give as many disciplines as you can whose aim is to study intelligent behavior of some sort. For each discipline, find out what aspect of behavior is investigated and what tools are used to study it. Be as liberal as you can regarding what defines intelligent behavior.

Exercise 1.3 Find out about two applications of AI (not classes of applications, but specific programs). For each application, write, at most, one typed page describing it. You should try to cover the following questions.

- (a) What does the application actually do (e.g., control a spacecraft, diagnose a photocopier, provide intelligent help for computer users)?
- (b) What AI technologies does it use (e.g., model-based diagnosis, belief networks, semantic networks, heuristic search, constraint satisfaction)?
- (c) How well does it perform? (According to the authors or to an independent review? How does it compare to humans? How do the authors know how well it works?)
- (d) Is it an experimental system or a fielded system? (How many users does it have? What expertise do these users require?)
- (e) Why is it intelligent? What aspects of it makes it an intelligent system?
- (f) [optional] What programming language and environment was it written in? What sort of user interface does it have?
- (g) References: Where did you get the information about the application? To what books, articles, or web pages should others who want to know about the application refer?

Exercise 1.4 For each of the Winograd schemas in Example 1.2, what knowledge is required to correctly answer the questions? Try to find a “cheap” method to find the answer, such as by comparing the number of results in a Google search for different cases. Try this for six other Winograd schemas of Davis [2015]. Try to construct an example of your own.

Exercise 1.5 Choose four pairs of dimensions that were not compared in Section 1.5.10 on page 31. For each pair, give one commonsense example of where the dimensions interact.