Table of Contents

Preface 1
Foreword 3
What is the point of this book? 3
AI today 4
Why not yet? AI to AGI 4
Bringing it all together 5

SECTION I: Are Super-Intelligent Machines In Your Future? 7

What's in Section I 8
The outline of the argument 8
Chapter 1: Could You Become a Computer? 11
Automating your brain 11
Faster and bigger 14
Brain in the basement 16
Backups and the passage of time 18
Distributed intelligence 19
Body swapping 20
Immortality 21
Ideas to consider 23
Chapter 2: What is Intelligence? 25
A Special Theory of Intelligence 26
The General Theory of Intelligence 27
An example 29
A few more details 31
Robots 32
AGI vs brain simulation 32
Chapter 3: Are Intelligent Machines Possible? 35
Computer horsepower 35
Computer software 39
Conclusion 42
Chapter 4: Are Intelligent Machines Inevitable? 45
A few scenarios 47
Will self-driving cars kill people? 48
Can AI be regulated? 49
Chapter 5: Won't AGI be Dangerous? 51
Doom and gloom 51
The Future of Life 53
Short-term issues 53
Longer term 54
The elephant in the room 56

Bugs and unintended consequences 56
Is there any good news? 57

SECTION II: What Is Intelligence? 59

What's in Section II 60
Chapter 6: Evolving Intelligence 61
The basics of evolution 61
Pros and cons of development by evolution 64
The evolution of intelligence 67
DNA 68
Evolution in civilization—memes 69
Evolution in computer hardware 70
Sex and the single CPU 71
Evolution in software 72
Chapter 7: Synapses, Brains, Transistors, and CPUs 75
Brains 75
Brainstem 76
The cerebellum 77
The neocortex 79
Computers 82
Neurons and synapses 85
Transistors 90
Chapter 8: Protozoans, Insects, and Computers 93
The black box control system 95
Reactions 97
A hypothetical protozoan 98
Ants 100
Basically smart 100
Chapter 9: "Th ablty to rcgnz mening frm prtl inpt" 103
Pattern recognition 104
Learning 107
Goals 109
Behavior sequences 112
Memory 113
Chapter 10: Sight, Sound, and Knowledge 119
Sound vs. sight, time vs. space 119
Content-addressable memory 122
Associative memory 125
Knowledge 126
Chapter 11: Modeling, Simulation, and Imagination 137
Simulation/imagination 142
Paying attention 147
Abstract reasoning 148
Acting on your imagination 149

Conclusion 151
Chapter 12: Free Will and Consciousness 153
Free will 154
Consciousness 156
The feeling objection: "How can a machine feel?" 161
The Chinese Room objection: "Where is the consciousness?" 163
The simulation objection: "Will it be *real* consciousness?" ... 164
Chapter 13: How Will Systems Act? 167
Sensation/perception, actions, and goals 167
Recognition and the knowledge store 168
Modeling the world 169
Imagining the world and choosing actions 170
Being conscious or happy or sad 170
Summation 171
Useful shortcuts 172

SECTION III: The Future Of Intelligent Machines 175

What's in Section III 176
Chapter 14: The Future of AI 177
Symbolic AI 178
Neural networks 180
An analogy 181
Why aren't we further along? 182
The future of AI and AGI 184
Chapter 15: Genius 187
IQ and testing 190
The IQ of a machine 192
Chapter 16: Asimov Revisited 195
Are "Laws of Robotics" necessary? 196
The simplest law 197
Curiosity: a basic drive 198
Unintended consequences 198
The power of laws 199
Nature vs. nurture 200
Some possible AGI laws 200
Rights for computers 201
Summary 202
Chapter 17: Beyond the Turing Test 203
Issues with the Turing Test 203
Proposed adjustments 204
Summary 205
Chapter 18: Will Computers Revolt? 207
Scenario 1: the peaceful-coexistence scenario 208

Scenario 2: the mad-man scenario 211
Scenario 3: the mad-machine scenario 211
Scenario 4: the mad-mankind scenario 214
Longer-term outcome: the end result 216
Conclusion 223
Afterword: Memoirs of a Computer 225
Acknowledgements 229
Glossary of terms and abbreviations 230
Index 235