

# AI Crash Course

A fun and hands-on introduction to reinforcement learning, deep learning, and artificial intelligence with Python

**Hadelin de Ponteves**

**Packt**

BIRMINGHAM - MUMBAI

# Table of Contents

<b>Preface</b>	<b>vii</b>
<b>Chapter 1: Welcome to the Robot World</b>	<b>1</b>
<b>Beginning the AI journey</b>	<b>1</b>
<b>Four different AI models</b>	<b>2</b>
The models in practice	3
Fundamentals	3
Thompson Sampling	3
Q-learning	3
Deep Q-learning	3
Deep convolutional Q-learning	4
<b>Where can learning AI take you?</b>	<b>4</b>
Energy	4
Healthcare	4
Transport and logistics	5
Education	5
Security	5
Employment	5
Smart homes and robots	5
Entertainment and happiness	6
Environment	6
Economy, business, and finance	6
<b>Summary</b>	<b>7</b>
<b>Chapter 2: Discover Your AI Toolkit</b>	<b>9</b>
<b>The GitHub page</b>	<b>9</b>
<b>Colaboratory</b>	<b>11</b>
<b>Summary</b>	<b>15</b>

<b>Chapter 3: Python Fundamentals – Learn How to Code in Python</b>	<b>17</b>
<b>Displaying text</b>	<b>18</b>
Exercise	18
<b>Variables and operations</b>	<b>19</b>
Exercise	20
<b>Lists and arrays</b>	<b>20</b>
Exercise	22
<b>if statements and conditions</b>	<b>22</b>
Exercise	23
<b>for and while loops</b>	<b>24</b>
Exercise	27
<b>Functions</b>	<b>27</b>
Exercise	28
<b>Classes and objects</b>	<b>29</b>
Exercise	31
<b>Summary</b>	<b>31</b>
<b>Chapter 4: AI Foundation Techniques</b>	<b>33</b>
<b>What is Reinforcement Learning?</b>	<b>33</b>
<b>The five principles of Reinforcement Learning</b>	<b>34</b>
Principle #1 – The input and output system	34
Principle #2 – The reward	35
Principle #3 – The AI environment	37
Principle #4 – The Markov decision process	37
Principle #5 – Training and inference	38
Training mode	38
Inference mode	39
<b>Summary</b>	<b>40</b>
<b>Chapter 5: Your First AI Model – Beware the Bandits!</b>	<b>41</b>
<b>The multi-armed bandit problem</b>	<b>41</b>
<b>The Thompson Sampling model</b>	<b>42</b>
Coding the model	43
Understanding the model	47
What is a distribution?	48
Tackling the MABP	52
The Thompson Sampling strategy in three steps	55
The final touch of shaping your Thompson Sampling intuition	56
Thompson Sampling against the standard model	57
<b>Summary</b>	<b>58</b>

---

<b>Chapter 6: AI for Sales and Advertising – Sell like the Wolf of AI Street</b>	<b>59</b>
<b>Problem to solve</b>	<b>59</b>
<b>Building the environment inside a simulation</b>	<b>61</b>
Running the simulation	64
Recap	66
<b>AI solution and intuition refresher</b>	<b>66</b>
AI solution	67
Intuition	68
<b>Implementation</b>	<b>68</b>
Thompson Sampling vs. Random Selection	69
Performance measure	69
Let's start coding	69
The final result	74
<b>Summary</b>	<b>76</b>
<b>Chapter 7: Welcome to Q-Learning</b>	<b>77</b>
<b>The Maze</b>	<b>78</b>
Beginnings	78
Building the environment	79
The states	79
The actions	80
The rewards	81
Building the AI	85
The Q-value	85
The temporal difference	86
The Bellman equation	87
Reinforcement intuition	88
<b>The whole Q-learning process</b>	<b>88</b>
Training mode	89
Inference mode	89
<b>Summary</b>	<b>90</b>
<b>Chapter 8: AI for Logistics – Robots in a Warehouse</b>	<b>91</b>
<b>Building the environment</b>	<b>94</b>
The states	94
The actions	95
The rewards	95
AI solution refresher	96
Initialization (first iteration)	96
Next iterations	96
<b>Implementation</b>	<b>97</b>
Part 1 – Building the environment	98
Part 2 – Building the AI Solution with Q-learning	101

---

Part 3 – Going into production	103
Improvement 1 – Automating reward attribution	105
Improvement 2 – Adding an intermediate goal	108
<b>Summary</b>	<b>111</b>
<b>Chapter 9: Going Pro with Artificial Brains – Deep Q-Learning</b>	<b>113</b>
<b>Predicting house prices</b>	<b>114</b>
Uploading the dataset	114
Importing libraries	116
Excluding variables	117
Data preparation	119
Scaling data	119
Building the neural network	122
Training the neural network	123
Displaying results	123
<b>Deep learning theory</b>	<b>125</b>
The neuron	125
Biological neurons	125
Artificial neurons	127
The activation function	128
The threshold activation function	129
The sigmoid activation function	130
The rectifier activation function	131
How do neural networks work?	133
How do neural networks learn?	135
Forward-propagation and back-propagation	136
Gradient Descent	137
Batch gradient descent	140
Stochastic gradient descent	143
Mini-batch gradient descent	145
<b>Deep Q-learning</b>	<b>145</b>
The Softmax method	147
Deep Q-learning recap	149
Experience replay	149
The whole deep Q-learning algorithm	150
<b>Summary</b>	<b>151</b>
<b>Chapter 10: AI for Autonomous Vehicles – Build a Self-Driving Car</b>	<b>153</b>
<b>Building the environment</b>	<b>153</b>
Defining the goal	156
Setting the parameters	160
The input states	163

---

The output actions	165
The rewards	166
<b>AI solution refresher</b>	<b>168</b>
<b>Implementation</b>	<b>169</b>
Step 1 – Importing the libraries	170
Step 2 – Creating the architecture of the neural network	171
Step 3 – Implementing experience replay	175
Step 4 – Implementing deep Q-learning	177
<b>The demo</b>	<b>188</b>
Installing Anaconda	189
Creating a virtual environment with Python 3.6	190
Installing PyTorch	192
Installing Kivy	194
<b>Summary</b>	<b>205</b>
<b>Chapter 11: AI for Business –</b>	
<b>Minimize Costs with Deep Q-Learning</b>	<b>207</b>
<b>Problem to solve</b>	<b>207</b>
<b>Building the environment</b>	<b>208</b>
Parameters and variables of the server environment	208
Assumptions of the server environment	209
Assumption 1 – We can approximate the server temperature	209
Assumption 2 – We can approximate the energy costs	210
Simulation	211
Overall functioning	212
Defining the states	214
Defining the actions	214
Defining the rewards	215
Final simulation example	216
<b>AI solution</b>	<b>220</b>
The brain	221
Implementation	223
Step 1 – Building the environment	224
Step 2 – Building the brain	232
Without dropout	233
With dropout	237
Step 3 – Implementing the deep reinforcement learning algorithm	238
Step 4: Training the AI	245
No early stopping	246
Early stopping	254
Step 5 – Testing the AI	256
<b>The demo</b>	<b>258</b>

<b>Recap – The general AI framework/Blueprint</b>	<b>268</b>
<b>Summary</b>	<b>270</b>
<b>Chapter 12: Deep Convolutional Q-Learning</b>	<b>271</b>
<b>What are CNNs used for?</b>	<b>271</b>
<b>How do CNNs work?</b>	<b>273</b>
Step 1 – Convolution	275
Step 2 – Max pooling	277
Step 3 – Flattening	280
Step 4 – Full connection	282
<b>Deep convolutional Q-learning</b>	<b>284</b>
<b>Summary</b>	<b>285</b>
<b>Chapter 13: AI for Games – Become the Master at Snake</b>	<b>287</b>
<b>Problem to solve</b>	<b>288</b>
<b>Building the environment</b>	<b>288</b>
Defining the states	289
Defining the actions	290
Defining the rewards	292
<b>AI solution</b>	<b>293</b>
The brain	293
The experience replay memory	295
<b>Implementation</b>	<b>295</b>
Step 1 – Building the environment	296
Step 2 – Building the brain	303
Step 3 – Building the experience replay memory	307
Step 4 – Training the AI	309
Step 5 – Testing the AI	315
<b>The demo</b>	<b>317</b>
Installation	317
The results	323
<b>Summary</b>	<b>325</b>
<b>Chapter 14: Recap and Conclusion</b>	<b>327</b>
<b>Recap – The general AI framework/blueprint</b>	<b>327</b>
<b>Exploring what's next for you in AI</b>	<b>329</b>
Practice, practice, and practice	329
Networking	330
Never stop learning	331
<b>Other Books You May Enjoy</b>	<b>333</b>
<b>Index</b>	<b>337</b>

---