

CONTENTS

1	Introduction	1
1	<i>Origins of AI</i>	3
2	<i>Narrow and General AI</i>	6
3	<i>Defining AI</i>	7
3.1	<i>Human-Centric Definitions</i>	9
3.2	<i>Rationalist Definitions</i>	13
3.3	<i>The Sceptics</i>	15
3.4	<i>Our Definition</i>	15
4	<i>AI, AI Everywhere</i>	22
5	<i>Superintelligence</i>	28
6	<i>Optimists, Pessimists and Pragmatists</i>	32
7	<i>If Not Now, When?</i>	34
8	<i>Robot Rules</i>	36
2	Unique Features of AI	39
1	<i>Sceptics of Novelty: Of Horses and HTTP</i>	39
2	<i>Fundamental Legal Concepts</i>	42
2.1	<i>Subjects and Agents</i>	42
2.2	<i>Causation</i>	57
3	<i>Features of AI Which Challenge Fundamental Legal Concepts</i>	64
3.1	<i>AI Makes Moral Choices</i>	64
3.2	<i>Independent Development</i>	70

3.3	<i>Why AI Is Not Like Chemicals or Biological Products</i>	78
4	<i>Conclusions on the Unique Features of AI</i>	79
3	Responsibility for AI	81
1	<i>Private and Criminal Law Distinguished</i>	81
2	<i>Private Law</i>	83
2.1	<i>Negligence</i>	84
2.2	<i>Strict and Product Liability</i>	91
2.3	<i>Vicarious Liability</i>	98
2.4	<i>No-Fault Accident Compensation Scheme</i>	102
2.5	<i>Contract</i>	106
2.6	<i>Insurance</i>	112
3	<i>Criminal Law</i>	117
3.1	<i>How Would Criminal Law Be Applied to Humans for the Actions of AI?</i>	118
3.2	<i>Advantages of Humans Being Criminally Responsible for AI</i>	119
3.3	<i>Shortcomings of Humans Being Held Criminally Responsible for AI</i>	120
4	<i>Responsibility for Beneficial Acts: AI and IP</i>	121
4.1	<i>Copyright</i>	122
4.2	<i>Case Study: The “Monkey Selfie” Case</i>	126
4.3	<i>Patents and Other Protections</i>	128
5	<i>Free Speech and Hate Speech</i>	129
6	<i>Conclusions on Responsibility for AI</i>	132
4	Rights for AI	133
1	<i>What Are Rights?</i>	134
1.1	<i>Hohfeld’s Incidents</i>	134
1.2	<i>Rights as Fictions</i>	135
2	<i>Animals: Man’s Best Friends?</i>	137
2.1	<i>Cultural Relativity of Attitudes to Animals</i>	137
2.2	<i>Animal Rights Through History</i>	139
3	<i>How the Human Got Her Rights</i>	143
4	<i>Why Robot Rights?</i>	145
4.1	<i>The Argument from Pain: “Suffer Little Robots”</i>	146
4.2	<i>The Argument from Compassion</i>	155
4.3	<i>The Argument from Value to Humanity</i>	164

4.4	<i>The Argument from Post-humanism: Hybrids, Cyborgs and Electronic Brains</i>	166
5	<i>Conclusions on Rights for AI</i>	170
5	Legal Personality for AI	173
1	<i>The Missing Link?</i>	173
2	<i>Is Legal Personality for AI Possible?</i>	175
2.1	<i>A Bundle of Rights and Obligations</i>	175
2.2	<i>Legal “Housing” of AI Within Existing Corporate Structures</i>	177
2.3	<i>New Legal Persons</i>	179
2.4	<i>Mutual Recognition of Foreign Legal Persons</i>	180
2.5	<i>Robots in the Boardroom</i>	181
3	<i>Should We Grant AI Legal Personality?</i>	183
3.1	<i>Pragmatic Justifications: Setting the Threshold</i>	184
3.2	<i>Filling the Accountability Gap</i>	185
3.3	<i>Encouraging Innovation and Economic Growth</i>	187
3.4	<i>Distributing the Fruits of AI Creativity</i>	188
3.5	<i>Skin in the Game</i>	188
3.6	<i>Arguments Against Personality for AI</i>	189
4	<i>Remaining Challenges</i>	197
4.1	<i>When Does AI Qualify for Legal Personality?</i>	197
4.2	<i>Identification of the AI</i>	197
4.3	<i>What Legal Rights and Responsibilities Might AI Hold?</i>	200
4.4	<i>Would Anyone Own AI?</i>	202
4.5	<i>Could a Robot Commit a Crime?</i>	202
5	<i>Conclusions on Legal Personality for AI</i>	205
6	Building a Regulator	207
1	<i>Why We Must Design Institutions Before We Can Write Laws</i>	207
1.1	<i>Philosophy of Institutional Design</i>	207
1.2	<i>AI Needs Principles Formulated by Public Bodies, Not Private Companies</i>	209
1.3	<i>Impartiality and Regulatory Capture</i>	211
1.4	<i>Too Many Rules, and Too Few</i>	212
2	<i>Rules for AI Should Be Made on a Cross-Industry Basis</i>	213

2.1	<i>The Shift from Narrow to General AI</i>	214
2.2	<i>The Need for General Principles</i>	218
3	<i>New Laws for AI Should Be Made by Legislation, Not Judges</i>	221
3.1	<i>Civil Law Systems</i>	221
3.2	<i>Common Law Systems</i>	221
3.3	<i>Writing Rules for AI</i>	222
4	<i>Current Trends in Government AI Regulation</i>	225
4.1	<i>UK</i>	225
4.2	<i>France</i>	228
4.3	<i>EU</i>	229
4.4	<i>USA</i>	230
4.5	<i>Japan</i>	231
4.6	<i>China</i>	232
4.7	<i>Conclusions on Current Trends in Government AI Regulation</i>	236
5	<i>International Regulation</i>	237
5.1	<i>An International Regulatory Agency for AI</i>	237
5.2	<i>Arbitrary Nature of National Boundaries</i>	237
5.3	<i>Cost of Uncertainty</i>	238
5.4	<i>Avoiding Arbitrage</i>	239
6	<i>Why Would Countries Agree to a Global Code?</i>	239
6.1	<i>Balancing Nationalism and Internationalism</i>	239
6.2	<i>Case Study: ICANN</i>	240
6.3	<i>Self-Interest and Altruism</i>	242
6.4	<i>Case Study: Space Law</i>	244
7	<i>Applying International Law to AI: The Toolbox</i>	247
7.1	<i>Traditional Structure of Public International Law</i>	247
7.2	<i>Subsidiarity</i>	249
7.3	<i>Varying Intensity of Regulation</i>	250
7.4	<i>Model Laws</i>	252
7.5	<i>An International Academy for AI Law and Regulation</i>	253
8	<i>Implementation and Enforcement of AI Laws</i>	254
8.1	<i>Coordination with National Regulators</i>	254
8.2	<i>Monitoring and Inspections</i>	256
8.3	<i>Sanctions for Non-compliance</i>	257
8.4	<i>Case Study: The EU's Sanctioning Method for Member States</i>	258

8.5	<i>Case Study: OECD Guidelines for Multinational Enterprises</i>	259
9	<i>Conclusions on Building a Regulator</i>	262
7	Controlling the Creators	263
1	<i>Creators and Creations</i>	263
2	<i>A Moral Regulator: The Quest for Legitimacy</i>	264
2.1	<i>“...of the People, by the People, for the People”</i>	265
2.2	<i>Case Study: GM Crops and Food Safety</i>	266
3	<i>Collaborative Lawmaking</i>	269
3.1	<i>Multidisciplinary Experts</i>	274
3.2	<i>Stakeholders, Interest Groups and NGOs</i>	274
3.3	<i>Companies</i>	275
3.4	<i>Case Study: The FCA FinTech Sandbox</i>	276
3.5	<i>Industry Standards Bodies</i>	279
3.6	<i>Case Study: The UK Human Fertilisation and Embryology Authority</i>	280
3.7	<i>A Minister for AI?</i>	282
4	<i>Proposed Regulatory Codes</i>	283
4.1	<i>The Roboethics Roadmap</i>	283
4.2	<i>The EPSRC and AHRC “Principles of Robotics”</i>	284
4.3	<i>CERNA Ethics of Robot Research</i>	285
4.4	<i>Asilomar 2017 Principles</i>	287
4.5	<i>IEEE Ethically Aligned Design</i>	291
4.6	<i>Microsoft Principles</i>	294
4.7	<i>EU Initiatives</i>	295
4.8	<i>Japanese Initiatives</i>	298
4.9	<i>Chinese Initiatives</i>	300
5	<i>Themes and Trends</i>	303
6	<i>Licensing and Education</i>	304
6.1	<i>Historic Guilds</i>	304
6.2	<i>Modern Professions</i>	305
6.3	<i>A Hippocratic Oath for AI Professionals</i>	307
6.4	<i>A Global Professional Body</i>	309
6.5	<i>AI Auditors</i>	310
6.6	<i>Objections and Responses</i>	311
7	<i>Regulating the Public: A Driver’s License for AI</i>	315

7.1	<i>Automatic for the People</i>	315
7.2	<i>How Might a Public AI License Function?</i>	316
8	<i>Conclusions on Controlling the Creators</i>	317
8	Controlling the Creations	319
1	<i>Laws of Identification</i>	320
1.1	<i>What Are Laws of Identification?</i>	320
1.2	<i>Why Might We Need Laws of Identification?</i>	321
1.3	<i>How Could Laws of Identification Be Achieved?</i>	323
2	<i>Laws of Explanation</i>	324
2.1	<i>What Are Laws of Explanation?</i>	324
2.2	<i>Why Might We Need Laws of Explanation?</i>	325
2.3	<i>How Could Laws of Explanation Be Achieved?</i>	325
3	<i>Laws on Bias</i>	336
3.1	<i>What Is Bias?</i>	337
3.2	<i>Why Might We Need Laws Against Bias?</i>	337
3.3	<i>How Could Laws Against Bias Be Achieved?</i>	345
4	<i>Limits on Limitation of AI Use</i>	350
4.1	<i>What are Laws of Limitation?</i>	350
4.2	<i>Why Might We Need Laws of Limitation?</i>	350
4.3	<i>How Could Laws of Limitation be Achieved?</i>	352
5	<i>The Kill Switch</i>	357
5.1	<i>What Is a Kill Switch?</i>	357
5.2	<i>Why Might We Need a Kill Switch?</i>	358
5.3	<i>How Could a Kill Switch Be Achieved?</i>	362
6	<i>Conclusions on Controlling the Creations</i>	369
9	Epilogue	371
	Index	373