

Contents

| | | | |
|--|----|---|-------------|
| Preface | | | viii |
| Chapter 1 What Is AI? | | | 1 |
| Chapter Overview | 1 | | 9 |
| 1.1 Introduction | 1 | 1.4 AI History | 9 |
| 1.2 Definitions of AI | 2 | Check Your Understanding | 13 |
| 1.2.1 Artificial Versus Natural Intelligence | 2 | 1.5 Major Types of AI for Business | 13 |
| 1.2.2 Other Areas of Natural Intelligence | 3 | Check Your Understanding | 16 |
| Check Your Understanding | 4 | 1.6 Ethical Considerations | 16 |
| 1.2.3 AI Versus Rule-Based Systems | 4 | Ethics in Action: ChatGPT Makes Things Up | 16 |
| Check Your Understanding | 7 | 1.7 Chapter Summary | 17 |
| 1.3 How Learning Algorithms Work: Label, Train, Predict | 7 | Key Terms | 18 |
| 1.3.1 Deep Learning Versus Classical Machine-Learning Algorithms | 8 | References | 21 |
| Chapter 2 Programming Tools for Deep Learning | | | 24 |
| Chapter Overview | 24 | 2.3.2 What Is a CLI? | 35 |
| 2.1 Python | 24 | 2.3.3 What Is Bash? | 35 |
| Check Your Understanding | 25 | 2.3.4 Ten Essential Bash Commands | 35 |
| 2.1.1 Custom Functions | 25 | 2.4 Key Packages for Deep Learning | 35 |
| Check Your Understanding | 28 | 2.4.1 Numerical Python—NumPy— https://numpy.org/ | 37 |
| 2.1.2 Class Inheritance | 28 | 2.4.1.1 NumPy for Data Representation | 37 |
| Check Your Understanding | 30 | 2.4.1.2 NumPy for Data Preprocessing | 38 |
| 2.1.3 Making Sense of Python Code and Graphs | 31 | 2.4.1.3 NumPy for Linear Algebra | 40 |
| 2.1.4 Common Coding Issues and Debugging Tips | 32 | Check Your Understanding | 41 |
| Check Your Understanding | 33 | 2.4.2 TensorFlow— https://www.tensorflow.org/ | 41 |
| 2.1.5 Python and Package Documentation— Why You Should Always Read It | 33 | 2.4.3 Keras— https://keras.io/ | 42 |
| 2.2 Google Colaboratory (Colab) | 34 | 2.4.4 Scikit-learn— https://scikit-learn.org/ | 43 |
| 2.3 The Terminal, Command Line Interface (CLI) and Bash | 35 | 2.5 Alternative AI Frameworks | 43 |
| 2.3.1 What Is the Terminal? | 35 | 2.6 Chapter Summary | 44 |
| | | Key Terms | 46 |
| Chapter 3 Mathematical Foundations of AI Models | | | 47 |
| Chapter Overview | 47 | 3.2.1.1 Central Tendency of a Numeric Feature | 54 |
| 3.1 Foundational Math for the Studies of AI | 47 | Ethics in Action: Are AI Models Trained to Be Dangerous? | 55 |
| 3.1.1 Basic Notations | 48 | Check Your Understanding | 55 |
| Check Your Understanding | 49 | 3.2.1.2 Spread of a Numeric Feature | 56 |
| 3.1.2 Exponents | 49 | Check Your Understanding | 58 |
| Check Your Understanding | 50 | 3.2.1.3 Standardization and Normalization | 58 |
| 3.1.3 Logarithm | 51 | Check Your Understanding | 60 |
| Check Your Understanding | 53 | 3.2.1.4 Skewness | 60 |
| 3.2 Probability and Statistics | 53 | Check Your Understanding | 61 |
| Why Does This Matter? | 53 | 3.2.2 Categorical Features | 61 |
| 3.2.1 Numeric Features | 54 | Check Your Understanding | 63 |

| | | | | | |
|---|---|-----|-------|--|------------|
| 3.3 | Linear Algebra | 64 | | Check Your Understanding | 74 |
| | Why Does This Matter? | 64 | 3.4 | Optimization | 74 |
| 3.3.1 | Data Representation | 64 | | Why Does This Matter? | 74 |
| | Ethics in Action: Oversimplified Reality | 67 | | Ethics in Action: Singular Focus on Cost Reduction | 75 |
| | Check Your Understanding | 68 | | Check Your Understanding | 75 |
| 3.3.2 | Vector Similarity | 68 | 3.5 | Calculus | 75 |
| 3.3.2.1 | Euclidean Distance | 68 | | Why Does This Matter? | 75 |
| | Check Your Understanding | 69 | | Check Your Understanding | 78 |
| 3.3.2.2 | Dot Product | 69 | 3.6 | Chapter Summary | 78 |
| | Check Your Understanding | 70 | | Key Terms | 80 |
| 3.3.2.3 | Cosine Similarity | 70 | | References | 82 |
| | Check Your Understanding | 71 | | | |
| 3.3.3 | Matrix Operations | 71 | | | |
| Chapter 4 Estimating Values | | | | | 83 |
| | Chapter Overview | 83 | 4.2.2 | Rectified Linear Unit (ReLU) | 96 |
| 4.1 | Introduction | 83 | | Check Your Understanding | 98 |
| 4.1.1 | Case in Point: Neural Zestimate as a Neural Network Model | 83 | 4.3 | Deep Learning for Regression Versus Linear Regression | 98 |
| 4.1.2 | Learning Algorithm Objectives: Regression Versus Classification | 84 | | Neural Network Playground by TensorFlow | 101 |
| 4.1.2.1 | Management and Operations | 85 | 4.4 | Lab: Configure a Sequential() Model for Regression | 101 |
| 4.1.2.2 | Finance and Accounting | 85 | 4.4.1 | Step 0: Load Libraries | 101 |
| 4.1.2.3 | Sales and Marketing | 85 | 4.4.2 | Step 1: Load and Prep Data Frame | 101 |
| | Check Your Understanding | 86 | 4.4.3 | Step 2: Prep Train Versus Test Sets | 102 |
| 4.1.3 | Learning Algorithm for Predicting Numeric Values | 86 | 4.4.4 | Step 3: Train Model | 102 |
| | Check Your Understanding | 88 | | Check Your Understanding | 104 |
| | Check Your Understanding | 93 | 4.4.5 | Step 4: Evaluate Model Performance | 106 |
| 4.2 | Deep Learning as a Multilayer Neural Network | 93 | | Ethics in Action: Transportation Disaster | 106 |
| | Check Your Understanding | 95 | | Check Your Understanding | 107 |
| 4.2.1 | Linear Activation Function | 96 | 4.5 | Chapter Summary | 107 |
| | | | | Key Terms | 109 |
| | | | | References | 110 |
| Chapter 5 Making Categorical Decisions | | | | | 111 |
| | Chapter Overview | 111 | | Check Your Understanding | 124 |
| 5.1 | Introduction | 111 | | Ethics in Action: When False Negatives Lead to Tragedies | 124 |
| | Check Your Understanding | 112 | 5.4 | One-Hot Encoding Lab | 125 |
| 5.1.1 | Case in Point: International Car Rentals | 112 | 5.4.1 | keras.utils.to_categorical() | 126 |
| 5.2 | Deep Learning for Classification | 112 | 5.4.2 | pd.get_dummies() | 126 |
| 5.2.1 | Step 1 (Label) | 113 | 5.5 | Configure a Sequential() Model for Classification Lab | 127 |
| | Check Your Understanding | 113 | 5.5.1 | Load Libraries | 127 |
| 5.2.1.1 | Dummy Coding | 113 | 5.5.2 | Load and Prep Data Frame | 128 |
| 5.2.1.2 | One-Hot Coding | 114 | 5.5.3 | Prep Train Versus Test Sets | 129 |
| | Check Your Understanding | 115 | 5.5.4 | Train Model | 129 |
| 5.2.2 | Step 2 (Train) | 115 | | Check Your Understanding | 130 |
| | Check Your Understanding | 116 | 5.5.5 | Evaluate Model Performance | 130 |
| 5.2.2.1 | Sigmoid for Binary Y Target | 116 | | Ethics in Action: AI as a Black Box | 133 |
| | Check Your Understanding | 118 | 5.6 | Chapter Summary | 133 |
| 5.2.2.2 | Softmax Function for Multiclass Y Target | 119 | | Key Terms | 135 |
| | Check Your Understanding | 122 | | References | 136 |
| | Ethics in Action: Only Men Qualify | 122 | | | |
| 5.2.3 | Step 3 (Predict) | 122 | | | |
| 5.3 | Model Evaluation | 123 | | | |

| | | |
|---|-----|---|
| Chapter 6 Evaluating and Improving Model Performance | | 137 |
| Chapter Overview | 137 | Check Your Understanding 149 |
| 6.1 Introduction | 137 | 6.4.2 Loss Functions for Classification Models 149 |
| 6.1.1 Case in Point: Preventing Fraud in the Electronic Payments Industry | 137 | Ethics in Action: Some Errors Are More Costly Than Others 152 |
| 6.2 Loss Functions Versus Metrics | 138 | Check Your Understanding 152 |
| Check Your Understanding | 138 | 6.5 The Problem With Perfect Accuracy 153 |
| 6.3 Loss Functions Versus Metrics for Regression Models | 139 | Check Your Understanding 153 |
| Check Your Understanding | 142 | 6.6 Model Improvement 153 |
| 6.4 Loss Functions Versus Metrics for Classification Models | 143 | 6.6.1 Model Architecture 154 |
| 6.4.1 Metrics for Classification Models | 143 | 6.6.2 Managing Model Overfit 154 |
| Ethics in Action: Who Gets to Pick a Decision Threshold? | 148 | Check Your Understanding 158 |
| | | 6.7 Chapter Summary 158 |
| | | Key Terms 160 |
| | | References 161 |
| Chapter 7 Inner Workings of Deep Learning (DL) Model Training | | 162 |
| Chapter Overview | 162 | 7.3.1.4 Problem: Vanishing and Exploding Gradient 176 |
| 7.1 Introduction | 162 | 7.3.1.5 Solution 1: Batch Normalization 177 |
| 7.2 Activation Functions | 163 | 7.3.1.6 Solution 2: Gradient Clipping 177 |
| 7.2.1 Threshold Function | 164 | Check Your Understanding 177 |
| Check Your Understanding | 165 | 7.3.1.7 Problem: Dying ReLU 177 |
| 7.2.2 Sigmoid Function | 165 | Check Your Understanding 178 |
| Check Your Understanding | 166 | 7.4 Optimizers for Keras Model Compilation 178 |
| 7.2.3 Hyperbolic Tangent (Tanh) Function | 166 | 7.4.1 Stochastic Gradient Descent (SGD) 179 |
| Check Your Understanding | 167 | Check Your Understanding 180 |
| 7.2.4 ReLU Function | 167 | 7.4.2 Adagrad 180 |
| Check Your Understanding | 168 | 7.4.3 Root Mean Square Propagation (RMSProp) 181 |
| 7.2.5 Activation Functions by Layer | 168 | Check Your Understanding 181 |
| Check Your Understanding | 169 | 7.4.4 Adaptive Moment Estimation (Adam) 182 |
| 7.3 How DL Model Training Works | 169 | Check Your Understanding 182 |
| 7.3.1 Gradient Descent | 169 | 7.5 Hyperparameter Tuning 183 |
| Check Your Understanding | 171 | 7.5.1 Learning Rate 183 |
| 7.3.2 Backpropagation | 171 | 7.5.2 Number of Training Steps (Number of Epochs) 184 |
| Check Your Understanding | 172 | 7.5.3 Batch Size 184 |
| 7.3.1.1 Learning Rate | 172 | Check Your Understanding 185 |
| Check Your Understanding | 174 | 7.6 Chapter Summary 185 |
| 7.3.1.2 Problem: Nonconvexity | 174 | Key Terms 187 |
| 7.3.1.3 Solution: Momentum | 175 | References 189 |
| Check Your Understanding | 176 | |
| Chapter 8 Visual Intelligence | | 190 |
| Chapter Overview | 190 | Check Your Understanding 202 |
| 8.1 Introduction | 190 | 8.3.2 Step 2: One-Hot Encode the Y Target 202 |
| 8.1.1 Case in Point: Visual Intelligence for Operational Efficiency | 191 | Check Your Understanding 203 |
| 8.2 Constructing a Dataset for Image Recognition | 191 | 8.3.3 Step 3: Configure CNN Layers 203 |
| 8.2.1 Representing Images Digitally as X Features | 192 | 8.3.3.1 Convolutional Layers 204 |
| Check Your Understanding | 197 | Check Your Understanding 207 |
| 8.2.2 Labeling Y Targets | 198 | 8.3.3.2 Pooling Layer 207 |
| Ethics in Action: Y Labels Have Worldviews | 198 | Check Your Understanding 210 |
| 8.3 Convolutional Neural Networks (CNN) | 198 | 8.3.4 Putting It All Together 211 |
| 8.3.1 Step 1: Feed the Image Input Dataset | 199 | Check Your Understanding 213 |
| | | Ethics in Action: Did You Call Me a Gorilla? 214 |
| | | 8.3.4.1 Tuning a CNN Model 214 |

| | | | | | |
|---|---|-----|----------|--|------------|
| 8.3.4.2 | Check Your Understanding | 216 | | Ethics in Action: Using AI to Identify Shoplifters | 221 |
| | Using Pretrained Models | 216 | | | |
| | Check Your Understanding | 219 | 8.5 | Chapter Summary | 221 |
| 8.4 | Business Applications | 220 | | Key Terms | 224 |
| | | | | References | 225 |
| Chapter 9 Language—Embeddings and LSTM | | | | | 227 |
| | Chapter Overview | 227 | 9.2.2.2 | Embeddings | 237 |
| 9.1 | Introduction | 227 | | Check Your Understanding | 244 |
| 9.1.1 | Case in Point: Financial Sentiment Analysis | 227 | 9.3 | DL Architectures for Language | 244 |
| 9.1.2 | Language Model Training Overview | 228 | 9.3.1 | Recurrence-Based Models | 245 |
| 9.2 | Featurizing Language Digitally | 229 | 9.3.1.1 | Challenges With RNN Models | 249 |
| 9.2.1 | Tokenization | 229 | | Check Your Understanding | 250 |
| 9.2.1.1 | Normalize Each Document | 230 | 9.3.1.2 | LSTM Networks | 250 |
| 9.2.1.2 | Split a Document Into Tokens | 230 | | Ethics in Action: Content Moderation or Censorship? | 257 |
| 9.2.1.3 | Map Tokens to Integer IDs | 231 | | Check Your Understanding | 259 |
| 9.2.1.4 | Standardize the Length of Each Document | 234 | 9.4 | Chapter Summary | 259 |
| | Check Your Understanding | 235 | | Key Terms | 261 |
| 9.2.2 | Document Representation | 235 | | References | 262 |
| 9.2.2.1 | Bag-of-Words (BoW) | 235 | | | |
| | Check Your Understanding | 236 | | | |
| Chapter 10 Large Language Models | | | | | 264 |
| | Chapter Overview | 264 | | Check Your Understanding | 285 |
| 10.1 | Introduction | 264 | 10.4.3 | Vibe Coding | 285 |
| 10.1.1 | Case in Point: LLMs in Financial Services at USAA and Vanguard | 265 | 10.5 | Limits of Transformer-Based Models | 286 |
| 10.2 | Transformer-Based Models | 265 | 10.5.1 | Hallucination | 287 |
| 10.2.1 | Encoder | 266 | | Ethics in Action: Hallucinations Can Cost You! | 287 |
| 10.2.2 | Simple Self-Attention Mechanism | 267 | 10.5.2 | Physical Understanding | 288 |
| | Check Your Understanding | 270 | 10.5.3 | Numeric Reasoning | 288 |
| 10.2.3 | Trainable Self-Attention Mechanism | 270 | 10.5.4 | Temporal Reasoning | 288 |
| 10.2.4 | Encoder-Only Models | 273 | 10.5.5 | Data Leakage | 290 |
| | Check Your Understanding | 275 | | Check Your Understanding | 290 |
| 10.2.5 | Encoder-Decoder and Decoder-Only Models | 276 | 10.5.6 | AI Slop | 290 |
| | Check Your Understanding | 277 | 10.6 | Model Fine-Tuning | 291 |
| 10.2.6 | Model Evaluation | 278 | | Check Your Understanding | 293 |
| | Check Your Understanding | 279 | 10.7 | Reasoning | 293 |
| | Ethics in Action: Do You Have Permission to Train an AI Model on My Data? | 280 | 10.8 | Context Engineering | 294 |
| 10.3 | Case Study: BloombergGPT | 280 | | Check Your Understanding | 295 |
| 10.3.1 | BloombergGPT Training Corpus | 280 | 10.9 | Agentic LLM Systems | 295 |
| 10.3.2 | Preprocessing and Tokenization for BloombergGPT Model Training | 280 | | Check Your Understanding | 296 |
| 10.3.3 | BloombergGPT Model Architecture | 280 | 10.9.1 | Create an LLM agent | 296 |
| 10.3.4 | BloombergGPT Model Training | 281 | 10.9.2 | LLM Agent With Custom-Made Tools | 297 |
| 10.3.5 | BloombergGPT Model Evaluation | 281 | 10.9.3 | LLM Agent With a Web Search Tool from an External MCP Server | 299 |
| 10.3.6 | BloombergGPT Use Cases | 281 | 10.9.4 | LLM Agent With Memory | 300 |
| 10.3.7 | BloombergGPT Implications | 282 | | Check Your Understanding | 302 |
| | Check Your Understanding | 282 | 10.9.5 | Additional Agentic Capabilities | 302 |
| 10.4 | Use of Pretrained LLMs | 282 | 10.9.5.1 | Reflection | 302 |
| 10.4.1 | Case Study: Google Gemini | 283 | 10.9.5.2 | Multiagent Collaboration | 303 |
| 10.4.2 | Prompt Engineering | 284 | | Ethics in Action: Potential Dangers of an AI Agent That Does Everything Autonomously | 304 |

| | | | | | |
|---|--|-----|----------|--|------------|
| 10.9.5.3 | Multimodality | 304 | 10.10 | Chapter Summary | 306 |
| 10.9.5.4 | Adaptation | 305 | | Key Terms | 309 |
| 10.9.5.5 | Safety Guardrails | 305 | | References | 313 |
| | Check Your Understanding | 306 | | | |
| Chapter 11 Recommenders | | | | | 319 |
| | Chapter Overview | 319 | 11.4 | Collaborative Filtering With DL Models | 331 |
| 11.1 | Introduction to Recommenders | 319 | 11.4.1 | Dimensionality Reduction With Matrix Factorization | 331 |
| 11.1.1 | Case in Point: Beauty Recommendation | 319 | 11.4.2 | Model-Based Recommendations With Keras Recommenders (Keras-RS) | 333 |
| 11.1.2 | Data Sources for Recommenders | 320 | 11.4.2.1 | Loading the Input Dataset | 333 |
| | Ethics in Action: Am I Constantly Being Watched? | 322 | 11.4.2.2 | Create Vocabularies | 335 |
| 11.2 | Content-Based Filtering | 322 | 11.4.2.3 | Preprocess the Data | 336 |
| | Check Your Understanding | 324 | 11.4.2.4 | Defining the RetrievalModel Class | 337 |
| | Ethics in Action: Recommenders Create Echo Chambers | 324 | 11.4.2.5 | Instantiating and Compiling the Model | 339 |
| 11.3 | Collaborative Filtering With Memory-Based Methods | 324 | 11.4.2.6 | Training the Keras-RS Model | 340 |
| 11.3.1 | Set-Based Methods | 325 | 11.4.2.7 | Ethics in Action: Netflix Prize | 341 |
| 11.3.1.1 | Jaccard Similarity | 325 | | Model Evaluation | 341 |
| 11.3.2 | Vector-Based Methods | 327 | | Check Your Understanding | 341 |
| 11.3.2.1 | Euclidean Distance, Dot Product, and Cosine Similarity | 327 | | Ethics in Action: Why Is This Recommended to Me? | 342 |
| 11.3.2.2 | Pearson Similarity | 329 | 11.5 | Chapter Summary | 342 |
| 11.3.3 | Choosing the “Right” Similarity Measure | 330 | | Key Terms | 344 |
| | Check Your Understanding | 331 | | References | 345 |
| Chapter 12 Ethics, Responsibility, and Risks | | | | | 347 |
| | Chapter Overview | 347 | 12.6.3 | AI Model Evaluation | 359 |
| 12.1 | Case Study: Can AI Run a Business? | 347 | 12.6.3.1 | Fairness by Calibration | 360 |
| 12.2 | AI Ethics | 348 | 12.6.3.2 | Fairness by Parity | 360 |
| 12.2.1 | What Is AI Ethics? | 348 | | Check Your Understanding | 361 |
| 12.3 | Principle 1: Beneficence | 349 | 12.7 | Principle 5: Explicability | 362 |
| 12.3.1 | Preference-Based Fine-Tuning | 349 | 12.7.1 | Explainability and Interpretability | 362 |
| 12.3.2 | Reinforcement Learning From Human Feedback (RLHF) | 350 | 12.7.1.1 | SHapley Additive exPlanations (SHAP) Values | 363 |
| 12.3.3 | Challenges to Alignment Assurance | 350 | 12.7.1.2 | LIME Values | 363 |
| | Check Your Understanding | 350 | 12.7.1.3 | Testing With Concept Activation Vectors (TCAV) | 364 |
| 12.4 | Principle 2: Nonmaleficence | 350 | | Check Your Understanding | 364 |
| 12.4.1 | Model Validation | 350 | 12.7.2 | Transparency and Traceability | 365 |
| 12.4.2 | Data Quality | 351 | 12.7.2.1 | Model Cards | 365 |
| 12.4.3 | Privacy and Anonymity | 352 | 12.7.2.2 | Algorithmic Audit | 366 |
| | Check Your Understanding | 355 | 12.7.3 | Legal and Professional Accountability | 366 |
| 12.4.4 | Autonomous Agents Present Additional Risks | 355 | 12.7.3.1 | Legal Accountability | 366 |
| 12.5 | Principle 3: Autonomy | 355 | 12.7.3.2 | Professional Accountability | 366 |
| 12.6 | Principle 4: Justice | 357 | 12.8 | Balancing Ethics and Business Interests | 367 |
| 12.6.1 | Data Collection and Preprocessing | 358 | | Check Your Understanding | 367 |
| 12.6.1.1 | Data Preprocessing | 358 | 12.9 | Chapter Summary | 367 |
| 12.6.1.2 | Protected (Sensitive) Variables | 358 | | Key Terms | 371 |
| 12.6.1.3 | Data Labeling | 358 | | References | 374 |
| 12.6.2 | AI Modeling | 359 | | | |
| Index of Key Terms | | | | | 380 |