An Intuitive Exploration of Artificial Intelligence

Simant Dube

An Intuitive Exploration of Artificial Intelligence

Theory and Applications of Deep Learning

Springer

Simant Dube
Technology and Innovation Office
Varian Medical Systems
A Siemens Healthineers Company
Palo Alto, CA, USA

ISBN 978-3-030-68623-9 ISBN 978-3-030-68624-6 (eBook)
https://doi.org/10.1007/978-3-030-68624-6

© The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer Nature Switzerland AG 2021
This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.
The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.
The publisher, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

"Learn how to see. Realize that everything connects to everything else."

Leonardo da Vinci

To my parents, for a warm home with lots of books,
to Mini, my life partner, for your unwavering encouragement and support, sine quo non,
to Yuvika and Saatvik, for the pitter-patter sound of your little feet every morning of your childhood,
to Sandra, for all the thought-provoking discussions about physics and mathematics with you,
to Saundarya, for all the chess, carrom, and badminton chill-out time with you,
to all lovers of science and mathematics endowed with natural *intelligence, which outperforms AI at the present,*
and to all the amazing AI machines in the future who will exceed my abilities in all respects despite their artificial *intelligence.*

Preface

> "I visualize a time when we will be to robots what dogs are to humans, and I'm rooting for the machines."
>
> Claude Shannon

> "Artificial intelligence would be the ultimate version of Google. The ultimate search engine that would understand everything on the web. It would understand exactly what you wanted, and it would give you the right thing. We're nowhere near doing that now. However, we can get incrementally closer to that, and that is basically what we work on."
>
> Larry Page

In June 2014, we drove to the Computer History Museum in Mountain View, California, not too far from where I was living. A spectacular surprise was in store for us. As soon as we entered the building, our eyes were riveted on an 11 ft. (3.4 m) long and 7 ft. (2.1 m) high sculpture of engineering weighing five tonnes and having 8000 parts. One hundred and sixty-five years after it was originally conceived and designed, the Difference Engine 2 of Charles Babbage was operational in front of us. I wonder if the Cambridge mathematician from Victorian England could foresee that one day his work would be on display in the heart of Silicon Valley. We listened with rapt attention to the docent talk and demonstration of this marvelous machine.

"I wish to God these calculations had been executed by steam!" an anguished Charles Babbage exclaims as he finds error after error in astronomy tables calculated by human hand. It is London and the summer of 1821, and this statement is a remarkably prescient wish for an era of automation by machines that would carry out intelligent tasks without any errors. After 200 years, Larry Page's vision of the

ultimate version of Google is not too different. Instead of the goal of automating tedious calculations, Google strives to build machines that will automatically understand everything on the World Wide Web.

What has changed in the intervening two centuries is the ubiquity of computing that has opened the doors to an unstoppable, colossal flood of data, also known as big data. Within the big data are numbers, observations, measurements, samples, metrics, pixels, voxels, 3-D range data, waveforms, words, symbols, fields, and records—which can all be processed and munged to expose structures and relationships. We want to use the big data to predict, classify, rank, decide, recommend, match, search, protect, analyze, visualize, understand, and reveal.

Of course, we are the creators of data, and we understand the underlying processes. We can bring our expertise to the fore and employ it to build machine learning models. This is the basis of classical ML, the parent of AI. Classical ML demands human labor during a process called feature engineering. It can work remarkably well for some applications, but it does have limits.

AI seeks to learn from the raw data. It chooses a large, deep neural network with tens or hundreds of millions of internal learnable parameters that are tweaked during training using calculus to minimize a suitably chosen loss function on a large annotated training dataset consisting of tens or hundreds of millions of examples. Once tweaking has been completed, the network can be deployed in production and used in the field to carry out the magical "intelligent" inference. During inference, AI models execute billions of multiplications, additions, and comparisons per second. If we were to create a mechanical slow-motion replica of such an AI model and witness a docent demonstration in a futuristic museum, it would truly be mind-boggling to see the countless moving parts of the endlessly sprawling machine as it crunched out a final answer.

This book is an endeavor to describe the science behind the making of such an impressive machine. One may ask—where is intelligence in the middle of billions of clinking, clanking parts? For human intelligence, we have first-hand subjective experience and we demand no further proof. For AI, the question is a difficult one. Let us take the first step of understanding how AI works.

Berkeley, CA, USA Simant Dube
January 2021

Contents

Part I Foundations

1 AI Sculpture 3
1.1 Manifolds in High Dimensions 4
1.2 Sculpting Process 6
1.3 Notational Convention 10
1.4 Regression and Classification 10
1.4.1 Linear Regression and Logistic Regression 11
1.4.2 Regression Loss and Cross-Entropy Loss 12
1.4.3 Sculpting with Shades 15
1.5 Discriminative and Generative AI 17
1.6 Success of Discriminative Methods 20
1.7 Feature Engineering in Classical ML 21
1.8 Supervised and Unsupervised AI 23
1.9 Beyond Manifolds 24
1.10 Chapter Summary 24

2 Make Me Learn 27
2.1 Learnable Parameters 28
2.1.1 The Power of a Single Neuron 28
2.1.2 Neurons Working Together 29
2.2 Backpropagation of Gradients 31
2.2.1 Partial Derivatives 32
2.2.2 Forward and Backward Passes 33
2.3 Stochastic Gradient Descent 37
2.3.1 Handling Difficult Landscapes 39
2.3.2 Stabilization of Training 40
2.4 Chapter Summary 42

3 Images and Sequences 45
3.1 Convolutional Neural Networks 46
3.1.1 The Biology of the Visual Cortex 46

3.1.2 Pattern Matching 47
3.1.3 3-D Convolution 48
3.2 Recurrent Neural Networks 51
3.2.1 Neurons with States 51
3.2.2 The Power of Recurrence 53
3.2.3 Going Both Ways 53
3.2.4 Attention 54
3.3 Self-Attention 57
3.4 LSTM 57
3.5 Beyond Images and Sequences 59
3.6 Chapter Summary 62

4 Why AI Works 65
4.1 Convex Polytopes 66
4.2 Piecewise Linear Function 67
4.2.1 Subdivision of the Input Space 67
4.2.2 Piecewise Non-linear Function 69
4.2.3 Carving Out the Feature Spaces 70
4.3 Expressive Power of AI 72
4.4 Convolutional Neural Network 74
4.5 Recurrent Neural Network 76
4.6 Architectural Variations 77
4.7 Attention and Carving 77
4.8 Optimization Landscape 81
4.8.1 Graph-Induced Polynomial 81
4.8.2 Gradient of the Loss Function 82
4.8.3 Visualization 83
4.8.4 Critical Points 84
4.9 The Mathematics of Loss Landscapes 87
4.9.1 Random Polynomial Perspective 87
4.9.2 Random Matrix Perspective 88
4.9.3 Spin Glass Perspective 91
4.9.4 Computational Complexity Perspective 92
4.9.5 SGD and Critical Points 93
4.9.6 Confluence of Perspectives 94
4.10 Distributed Representation and Intrinsic Dimension 95
4.11 Chapter Summary 97

5 Novice to Maestro 101
5.1 How AI Learns to Sculpt 102
5.1.1 Training Data 102
5.1.2 Evaluation Metrics 103
5.1.3 Hyperparameter Search 105
5.1.4 Regularization 105
5.1.5 Bias and Variance 107
5.1.6 A Fairy Tale in the Land of ML 108

5.2 Learning Curves 110
5.3 From the Lab to the Dirty Field 113
5.4 System Design 114
5.5 Flavors of Learning 116
5.6 Ingenuity and Big Data in the Success of AI 119
5.7 Chapter Summary 119

6 Unleashing the Power of Generation 121
6.1 Creating Universes 122
6.2 To Recognize It, Learn to Draw It 123
6.3 General Definition 124
6.4 Generative Parameters 125
6.5 Generative AI Models 126
6.5.1 Restricted Boltzmann Machines 126
6.5.2 Autoencoders 128
6.5.3 Variational Autoencoder 129
6.5.4 Pixel Recursive Models 131
6.5.5 Generative Adversarial Networks 131
6.5.6 Wasserstein Generative Adversarial Networks 132
6.6 The Carving Process in Generative AI 134
6.7 Representation of Individual Signals 136
6.8 Chapter Summary 137

7 The Road Most Rewarded 141
7.1 Reinforcement Learning 141
7.2 Learning an Optimal Policy 143
7.3 Deep Q-Learning 147
7.4 Policy Gradient 148
7.4.1 Intuition 149
7.4.2 Mathematical Analysis 151
7.5 Let's Play and Explore 153
7.6 Chapter Summary 156

8 The Classical World 159
8.1 Maximum Likelihood Estimation 160
8.2 Uncertainty in Estimation 161
8.3 Linear Models 162
8.3.1 Linear Regression 162
8.3.2 The Geometry of Linear Regression 163
8.3.3 Regularization 163
8.3.4 Logistic Regression 164
8.4 Classical ML 165
8.4.1 k-Nearest Neighbors 165
8.4.2 Naive Bayes Classifier 165
8.4.3 FDA, LDA, and QDA 165
8.4.4 Support Vector Machines 166

8.4.5 Neural Networks 166
8.4.6 Decision Trees 167
8.4.7 Gaussian Process Regression 167
8.4.8 Unsupervised Methods 168
8.5 XGBoost 169
8.5.1 Relevance Ranking 171
8.6 Chapter Summary 172

Part II Applications

9 To See Is to Believe 177
9.1 Image Classification 178
9.1.1 Motivating Examples 178
9.1.2 LeNet 179
9.1.3 Stacked Autoencoders 180
9.1.4 AlexNet 180
9.1.5 VGG 181
9.1.6 ResNet 181
9.1.7 Inception V3 182
9.1.8 Showroom of Models 182
9.1.9 Your Own Network 184
9.2 Object Detection as Classification 184
9.2.1 Sliding Window Method 186
9.2.2 Region Proposal Method 187
9.3 Regression on Images 188
9.3.1 Motivating Examples 188
9.3.2 Object Detection as Regression 188
9.3.2.1 Regression Output 189
9.3.2.2 Grid-Based Approach 189
9.4 Attention in Computer Vision 191
9.5 Semantic Segmentation 192
9.6 Image Similarity 195
9.7 Video Analysis 197
9.8 3-D Data 199
9.9 Self-Driving Vehicles 201
9.10 Present and Future 204
9.11 Protein Folding 205
9.12 Chapter Summary 206

10 Read, Read, Read 207
10.1 Natural Language Understanding 208
10.2 Embedding Words in a Semantic Space 209
10.3 Sequence to Sequence 211
10.3.1 Encoder-Decoder Architecture 212
10.3.2 Neural Machine Translation 212
10.4 Attention Mechanism 213

10.5 Self-Attention 214
10.6 Creativity in NLU Solutions 218
10.7 AI and Human Culture 222
10.8 Recommender Systems 224
10.9 Reward-Based Formulations 226
10.10 Chapter Summary 226

11 Lend Me Your Ear 227
11.1 Classical Speech Recognition 229
11.2 Spectrogram to Transcription 229
11.2.1 Alignment-Free Temporal Connections 229
11.2.2 End-to-End Solution 231
11.2.3 Don't Listen to Others 232
11.3 Speech Synthesis 233
11.4 Handwriting Recognition 235
11.5 Chapter Summary 236

12 Create Your Shire and Rivendell 237
12.1 From Neurons to Art 238
12.1.1 DeepDream 238
12.1.2 Style Transfer 239
12.2 Image Translation 243
12.3 DeepFake 244
12.4 Creative Applications 246
12.5 Chapter Summary 246

13 Math to Code to Petaflops 249
13.1 Software Frameworks 250
13.1.1 The Twentieth Century 250
13.1.2 The Twenty-First Century 250
13.1.3 AI Frameworks 251
13.2 Let's Crunch Numbers 253
13.2.1 Computing Hardware 254
13.2.2 GPU Machines 254
13.2.3 Cloud GPU Instances 255
13.2.4 Training Script 256
13.2.5 Deployment 256
13.3 Speeding Up Training 257
13.3.1 Data Parallelism 257
13.3.2 Delayed and Compressed SGD 259
13.4 Open Ecosystem and Efficient Hardware 259
13.5 Chapter Summary 260

14 AI and Business 261
14.1 Strategy 261
14.2 Organization 262
14.3 Execution 262

14.4 Evaluation 264
14.5 Startups 265
14.6 Chapter Summary 266

Part III The Road Ahead

15 Keep Marching On 269
15.1 Robust AI 270
15.1.1 Adversarial Examples 270
15.1.2 Learning from Human Vision 272
15.1.3 Fusion of Evidence 275
15.1.4 Interpretable AI 277
15.2 AI Extraordinaire 279
15.3 Chapter Summary 282

16 Benevolent AI for All 285
16.1 Benefits of AI 286
16.2 AI in Medicine 286
16.3 Dangers of AI 287
16.4 AI-Human Conflict 288
16.5 Choices Ahead 289
16.6 Chapter Summary 290

17 Am I Looking at Myself? 293
17.1 Is It Computable or Non-computable? 294
17.2 Is Consciousness Everywhere? 296
17.3 Who Is the Storyteller? 297
17.4 Chapter Summary 298

A Solutions 301

B Lab Exercises and Projects 323
B.1 Exercises 323
B.2 Exploration 325
B.3 Debate and Discussion 326

Further Reading 327

Glossary 329

References 335

Index 347

Acronyms

L^p	Lebesgue Space.
AGI	Artificial General Intelligence.
AI	Artificial Intelligence.
ALS	Alternating Least Squares.
ASIC	Application-Specific Integrated Circuit.
ASR	Automatic Speech Recognition.
AUC	Area Under Curve.
BERT	Bidirectional Encoder Representation from Transformer.
Capsnet	Capsule Neural Network.
CART	Classification and Regression Trees.
CNN	Convolutional Neural Network.
CPU	Central Processing Unit.
CTC	Connectionist Temporal Classification.
DBN	Deep Belief Network.
DBSCAN	Density-Based Spatial Clustering of Applications with Noise.
DNC	Differentiable Neural Computer.
DNN	Deep Neural Network.
DQN	Deep Q-Network.
DRL	Deep Reinforcement Learning.
E2E	End-to-End.
EM	Expectation Maximization.
FCN	Fully Convolutional Network.
FDA	Fisher Discriminant Analysis.
FFN	Feed-Forward Network.
FPGA	Field-Programmable Gate Array.
FPR	False Positive Rate.
GAN	Generative Adversarial Network.
GBM	Gradient Boosted Machine.
GPR	Gaussian Process Regression.
GPT	Generative Pretrained Transformer.
GPU	Graphics Processing Unit.

GRU Gated Recurrent Unit.
HMM Hidden Markov Model.
HOG Histogram of Oriented Gradients.
IOU Intersection Over Union.
JS Jensen-Shannon.
k-NN k-Nearest Neighbor.
KL Kullback-Leibler.
LAD Least Absolute Deviation.
LDA Linear Discriminant Analysis.
LSTM Long Short-Term Memory.
MAP Maximum a Posteriori.
mAP Mean Average Precision.
MART Multiple Additive Regression Trees.
MCMC Markov Chain Monte Carlo.
MCTS Monte Carlo Tree Search.
MDP Markov Decision Process.
MDS Multidimensional Scaling.
MFCC Mel Frequency Cepstral Coefficients.
ML Machine Learning.
MLE Maximum Likelihood Estimation.
MLP Multi-Layer Perceptron.
MOE Mixture of Experts.
MSE Mean Squared Error.
NAS Neural Architecture Search.
NCF Neural Collaborative Filtering.
NDCG Normalized Discounted Cumulative Gain.
NER Named-Entity Recognition.
NLP Natural Language Processing.
NLU Natural Language Understanding.
NMS Non-Maximum Suppression.
NMT Neural Machine Translation.
NTM Neural Turing Machine.
OCR Optical Character Recognition.
OLS Ordinary Least Squares.
PCA Principal Component Analysis.
PGM Probabilistic Graphical Model.
PLS Partial Least Squares.
PN Policy Network.
PR Precision Recall.
Q Quality (in reinforcement learning).
QA Question And Answer.
QDA Quadratic Discriminant Analysis.
R-CNN Region-based CNN.
RANSAC Random Sample Consensus.
RBM Restricted Boltzmann Machine.

ReLU	Rectified Linear Unit.
RL	Reinforcement Learning.
RNN	Recurrent Neural Network.
ROC	Receiver Operating Characteristic.
SARSA	State, Action, Reward, State, Action.
Seq2Seq	Sequence-to-Sequence.
SGD	Stochastic Gradient Descent.
SIFT	Scale-Invariant Feature Transform.
SOFM	Self-Organizing Feature Maps.
SRL	Statistical Relational Learning.
SSD	Single-Shot Detector.
SVD	Singular-Value Decomposition.
SVM	Support Vector Machine.
t-SNE	t-Stochastic Neighbor Embedding.
TD	Temporal Difference.
TPR	True Positive Rate.
TPU	Tensor Processing Unit.
VAE	Variational Autoencoder.
VQ	Vector Quantization.
XGBoost	Extreme Gradient Boosting.
YOLO	You Only Look Once.

Author's Note

This book is about developing conceptual understanding of Artificial Intelligence (AI), a rapidly moving field of growing complexity and of immense transformative potential. Deep learning has emerged as a key enabling component of AI. Significant attention is being paid to it by the media, governments, businesses, and the public. It can safely be said that AI will continue to remain an exciting frontier for humanity to seek and explore for centuries to come. Besides the motivation of using AI for its myriad practical applications, we humans like to contemplate and comprehend. This book seeks an "understanding" of AI and Machine Learning (ML) in the truest sense of the word. It is an earnest endeavor to unravel what is happening at the algorithmic level, grasp how applications are being built, and show the long adventurous road in the future. It is my ardent wish to answer some fundamental questions. Can we really create intelligent machines? Can an algorithm serve as a model of learning and intelligence? Will there always be some differences between human intelligence and AI? In order to answer these questions, the first step is to understand the state of the art in AI.

The book rests on a vast knowledge store in the form of articles, books, interviews, news stories, lectures, and software, which are the result of diligent arduous work by many scientists, mathematicians, engineers, educators, students, authors, and journalists. I acknowledge their contributions, as the book is really a distillation of their work. While writing the book, I felt like a small child standing next to all the great minds and watching the beckoning waves of future breakthroughs while they break against the present shore.

The primary goal of the book is to serve as a textbook for AI and deep learning at college and university level for students who have basic familiarity with calculus, linear algebra, and probability and statistics. When being used as a textbook, the book follows inverting-the-classroom philosophy, which states that learning is an active process of discovery, research, and contemplation. The student is encouraged to look up references in the bibliography to do exercises and to dive deeper into topics. An ideal course in AI will use it as a textbook and complement it with course projects and paper-reading sessions. The book emphasizes deep learning because it

has revolutionized the field of AI. When combined with other techniques such as reinforcement learning, it can lead to amazing results.

At the same time, the book is intended for anyone who wishes to learn about AI. Learning is lifelong, and I have met many people who want to keep themselves continually challenged by learning new fields. Mathematical notation is simplified, and various illustrations are used to convey the intuition. Because the target audience is anyone who wants to learn the fundamentals of AI, I have chosen geometrical insight and intuitive exposition as the guiding principles. Therefore, the writing style is often informal and based on analogies. The language of mathematics is employed to convey concepts with greater clarity when it is needed. The book emphasizes both the theoretical and applied sides of AI and, therefore, will be able to meet different learning needs.

This book project is a by-product of my own professional journey in computer science as the field grew, morphed, and changed over time. It is a personal journey of discovery and learning and of being gently nudged by my life partner, Mini, to share with the wider world the intuition behind this very fine application of mathematics and engineering. I hope you will enjoy the book.

Berkeley, CA, USA Simant Dube
January 2021