



Machine Learning and AI for Healthcare

Big Data for Improved Health Outcomes

Arjun Panesar

apress®

Machine Learning and AI for Healthcare

**Big Data for Improved Health
Outcomes**

Arjun Panesar

Apress®

Machine Learning and AI for Healthcare

Arjun Panesar
Coventry, UK

ISBN-13 (pbk): 978-1-4842-3798-4

ISBN-13 (electronic): 978-1-4842-3799-1

<https://doi.org/10.1007/978-1-4842-3799-1>

Library of Congress Control Number: 2018967454

Copyright © 2019 by Arjun Panesar

This work is subject to copyright. All rights are reserved 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.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr

Acquisitions Editor: Celestin Suresh John

Development Editor: Matthew Moodie

Coordinating Editor: Divya Modi

Cover designed by eStudioCalamar

Cover image designed by Freepik (www.freepik.com)

Book imagery designed by Krystal Sidwell

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail rights@apress.com, or visit <http://www.apress.com/rights-permissions>.

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at <http://www.apress.com/bulk-sales>.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub via the book's product page, located at www.apress.com/978-1-4842-3798-4. For more detailed information, please visit <http://www.apress.com/source-code>.

Printed on acid-free paper

*Dedicated to the giants on whose shoulders we stand.
They laid the foundations with hard work, determination,
blood, sweat, sacrifice, and tears. Tribute to the
foremothers and forefathers; Kirpa, Ananta.*

Table of Contents

About the Authorxix

About the Technical Reviewersxxi

Acknowledgmentsxxiii

Introductionxxv

Chapter 1: What Is Artificial Intelligence?1

 A Multifaceted Discipline1

 Examining Artificial Intelligence.....4

 Reactive Machines6

 Limited Memory—Systems That Think and Act Rationally6

 Theory of Mind—Systems That Think Like Humans6

 Self-Aware AI—Systems That Are Humans.....7

 What Is Machine Learning?.....8

 What Is Data Science?9

 Learning from Real-Time, Big Data10

 Applications of AI in Healthcare12

 Prediction13

 Diagnosis.....13

 Personalized Treatment and Behavior Modification13

 Drug Discovery14

 Follow-Up Care14

TABLE OF CONTENTS

Realizing the Potential of AI in Healthcare 15

 Understanding Gap 15

 Fragmented Data 15

 Appropriate Security 16

 Data Governance 16

 Bias 17

Software 17

Conclusion 18

Chapter 2: Data 21

 What Is Data? 21

 Types of Data 23

 Big Data 26

 Volume 28

 Variety 31

 Velocity 34

 Value 37

 Veracity 39

 Validity 41

 Variability 41

 Visualization 42

 Small Data 42

 Metadata 43

 Healthcare Data—Little and Big Use Cases 44

 Predicting Waiting Times 44

 Reducing Readmissions 44

 Predictive Analytics 45

 Electronic Health Records 45

TABLE OF CONTENTS

Value-Based Care/Engagement.....	46
Healthcare IoT—Real-Time Notifications, Alerts, Automation	47
Movement Toward Evidence-Based Medicine	49
Public Health	50
Evolution of Data and Its Analytics.....	51
Turning Data into Information: Using Big Data	53
Descriptive Analytics	54
Diagnostic Analytics	55
Predictive Analytics	55
Prescriptive Analytics	58
Reasoning	59
Deduction	60
Induction.....	60
Abduction	61
How Much Data Do I Need for My Project?	61
Challenges of Big Data.....	62
Data Growth.....	62
Infrastructure.....	62
Expertise.....	63
Data Sources	63
Quality of Data	63
Security	63
Resistance	64
Policies and Governance.....	65
Fragmentation.....	65
Lack of Data Strategy.....	65
Visualization.....	66

TABLE OF CONTENTS

Timeliness of Analysis.....	66
Ethics	66
Data and Information Governance	66
Data Stewardship	67
Data Quality	68
Data Security	68
Data Availability	68
Data Content.....	69
Master Data Management (MDM).....	69
Use Cases	69
Deploying a Big Data Project.....	71
Big Data Tools	72
Conclusion	73
Chapter 3: What Is Machine Learning?.....	75
Basics	77
Agent	77
Autonomy	78
Interface	78
Performance	79
Goals.....	79
Utility	79
Knowledge.....	80
Environment	80
Training Data	81
Target Function.....	82
Hypothesis.....	82
Learner	82

Hypothesis	82
Validation	82
Dataset	82
Feature	82
Feature Selection	83
What Is Machine Learning?.....	83
How Is Machine Learning Different from Traditional Software Engineering?	84
Machine Learning Basics.....	85
Supervised Learning.....	86
How Machine Learning Algorithms Work	95
How to Perform Machine Learning	96
Specifying the Problem	97
Preparing the Data.....	99
Choosing the Learning Method.....	102
Applying the Learning Methods.....	103
Assessing the Method and Results	107
Optimization	113
Reporting the Results	116
Chapter 4: Machine Learning Algorithms	119
Defining Your Machine Learning Project.....	120
Task (T)	120
Performance (P).....	121
Experience (E).....	121
Common Libraries for Machine Learning.....	123
Supervised Learning Algorithms	125
Classification	127
Regression.....	128

TABLE OF CONTENTS

Decision trees 129

 Iterative Dichotomizer 3 (ID3) 133

 C4.5 134

 CART 134

Ensembles..... 135

 Bagging 135

 Boosting 137

Linear Regression 139

Logistic Regression..... 141

SVM..... 143

Naive Bayes 145

kNN: k-nearest neighbor 147

Neural Networks 148

 Perceptron..... 149

 Artificial Neural Networks 151

Deep Learning..... 152

 Feedforward Neural Network 154

 Recurrent Neural Network (RNN)—Long Short-Term Memory..... 154

 Convolutional Neural Network..... 155

 Modular Neural Network 155

 Radial Basis Neural Network..... 156

Unsupervised Learning 157

 Clustering 158

 K-Means 158

 Association 160

 Apriori 161

Dimensionality Reduction Algorithms 162

Dimension Reduction Techniques	165
Missing/Null Values	165
Low Variance	165
High Correlation.....	165
Random Forest Decision Trees	166
Backward Feature Elimination.....	166
Forward Feature Construction	166
Principal Component Analysis (PCA).....	166
Natural Language Processing (NLP).....	167
Getting Started with NLP.....	170
Preprocessing: Lexical Analysis	170
Noise Removal.....	171
Lexicon Normalization	171
Porter Stemmer	171
Object Standardization	172
Syntactic Analysis	172
Dependency Parsing.....	173
Part of Speech Tagging.....	173
Semantic analysis.....	175
Techniques Used Within NLP	175
N-grams	175
TF IDF Vectors.....	176
Latent Semantic Analysis	177
Cosine Similarity.....	177
Naïve Bayesian Classifier	178
Genetic Algorithms.....	179

TABLE OF CONTENTS

Best Practices and Considerations	180
Good Data Management	180
Establish a Performance Baseline	181
Spend Time Cleaning Your Data	181
Training Time	182
Choosing an Appropriate Model	182
Choosing Appropriate Variables	182
Redundancy	183
Overfitting	183
Productivity	183
Understandability	184
Accuracy	184
Impact of False Negatives	184
Linearity	185
Parameters	185
Ensembles	186
Use Case: Type 2 Diabetes	186
Chapter 5: Evaluating Learning for Intelligence	189
Model Development and Workflow	190
Why Are There Two Approaches to Evaluating a Model?	191
Evaluation Metrics	192
Skewed Datasets, Anomalies, and Rare Data	199
Parameters and Hyperparameters	199
Tuning Hyperparameters	200
Hyperparameter Tuning Algorithms	200
Grid Search	201
Random Search	201

Multivariate Testing.....202

 Which Metric Should I Use for Evaluation?.....202

 Correlation Does Not Equal Causation.....203

 What Amount of Change Counts as Real Change?203

 Types of Tests, Statistical Power, and Effect Size204

 Checking the Distribution of Your Metric204

 Determining the Appropriate *p* Value.....204

 How Many Observations Are Required?205

 How Long to Run a Multivariate Test?205

 Data Variance206

 Spotting Distribution Drift.....206

 Keep a Note of Model Changes206

Chapter 6: Ethics of Intelligence.....207

 What Is Ethics?210

 What Is Data Science Ethics?210

 Data Ethics210

 Informed Consent.....212

 Freedom of Choice212

 Should a Person’s Data Consent Ever Be Overturned?213

 Public Understanding214

 Who Owns the Data?.....215

 What Can the Data Be Used For?218

 Privacy: Who Can See My Data?.....220

 How Will Data Affect the Future?221

 Prioritizing Treatments.....221

 Determining New Treatments and Management Pathways.....222

TABLE OF CONTENTS

More real-world evidence	222
Enhancements in Pharmacology	222
Optimizing Pathways Through Connectivity—Is There a Limit?	223
Security	223
Ethics of Artificial Intelligence and Machine Learning	224
Machine Bias	225
Data Bias	226
Human Bias	226
Intelligence Bias	226
Bias Correction	227
Is Bias a Bad Thing?	228
Prediction Ethics	228
Explaining Predictions	229
Protecting Against Mistakes	230
Validity	231
Preventing Algorithms from Becoming Immoral	231
Unintended Consequences	233
How Does Humanity Stay in Control of a Complex and Intelligent System?	234
Intelligence	235
Health Intelligence	237
Who Is Liable?	238
First-Time Problems	240
Defining Fairness	241
How Do Machines Affect Our Behavior and Interaction	241
Humanity	241
Behavior and Addictions	242
Economy and Employment	243

Affecting the future	244
Playing God	244
Overhype and Scaremongering.....	245
Stakeholder Buy-In and Alignment	245
Policy, Law, and Regulation.....	245
Data and Information Governance	246
Is There Such a Thing as Too Much Policy?	247
Global standards and schemas.....	247
Do We Need to Treat AI with Humanity?	248
Employing Data Ethics Within Your Organization	249
Ethical Code.....	249
Ethical Framework Considerations.....	251
A Hippocratic Oath for Data Scientists.....	253
Auditing Your Frameworks.....	253
Chapter 7: Future of Healthcare.....	255
Shifting from Volume to Value.....	256
Evidence-Based Medicine.....	261
Personalized Medicine	264
Vision of the Future.....	266
Connected Medicine	269
Disease and Condition Management.....	274
Virtual Assistants.....	275
Remote Monitoring	276
Medication Adherence	277
Accessible Diagnostic Tests.....	277
Smart Implantables.....	278

TABLE OF CONTENTS

Digital Health and Therapeutics278

Education279

Incentivized Wellness.....280

AI.....281

 Mining Records.....281

 Conversational AI.....282

 Making Better Doctors.....283

Virtual and Augmented Reality.....290

 Virtual Reality290

 Augmented Reality290

 Merged Reality291

 Pain Management.....291

 Physical Therapy.....292

 Cognitive Rehabilitation.....292

 Nursing and Delivery of Medicine292

 Virtual Appointments and Classrooms.....293

Blockchain294

 Verifying the Supply Chain296

 Incentivized Wellness296

 Patient Record Access.....297

Robots.....298

 Robot-Assisted Surgery.....298

 Exoskeletons298

 Inpatient Care299

 Companions.....299

 Drones299

Smart Places.....	300
Smart Homes.....	301
Smart Hospitals	302
Reductionism	303
Innovation vs. Deliberation.....	303
Chapter 8: Case Studies.....	305
Case Study Selection	305
Conclusion	307
Case Study: AI for Imaging of Diabetic Foot Concerns and Prioritization of Referral for Improvements in Morbidity and Mortality	307
Background	307
Cognitive Vision	309
Project Aims	310
Challenges.....	312
Conclusions	315
Case Study: Outcomes of a Digitally Delivered, Low-Carbohydrate, Type 2 Diabetes Self-Management Program: 1-Year Results of a Single-Arm Longitudinal Study	316
Background	316
Objectives.....	317
Methods.....	317
Results.....	319
Observations.....	319
Conclusions	320
Case Study: Delivering A Scalable and Engaging Digital Therapy for Epilepsy	321
Background	321
Implementing the Evidence Base	321

TABLE OF CONTENTS

Sensor-Driven Digital Program322

Research.....323

Project Impact324

Preliminary Analysis324

Case Study: Improving Learning Outcomes For Junior Doctors
Through the Novel Use of Augmented and Virtual Reality..... 325

 Background 325

 Aims..... 326

 Project Description 326

 Conclusions 327

Case Study: Big Data, Big Impact, Big Ethics: Diagnosing Disease
Risk from Patient Data 328

 Background 328

 Platform Services 329

 Medication Adherence, Efficacy and Burden 329

 Community Forum 330

 AI prioritization of patient interactions 331

 Real-World Evidence 332

 Ethical Implications of Predictive Analytics 333

 Integration of the IoT 334

 Conclusions 334

Technical Glossary335

References343

Index359

About the Author



Arjun Panesar is the founder of Diabetes Digital Media (DDM), who operate the world's largest diabetes community and provide evidence-based digital solutions. Arjun holds a first-class honors degree (MEng) in Computing and Artificial Intelligence from Imperial College, London. Benefiting from a decade of experience in big data and affecting user outcomes, Arjun leads the development of intelligent, evidence-based, digital health solutions that harness the power of big data and machine learning to provide precision healthcare to patients, health agencies, and governments worldwide.

Arjun's work has won numerous awards and international recognition, featured in the BBC, Forbes, Daily Mail, The Times and ITV.

Arjun is an advisor to the Information School, University of Sheffield.

About the Technical Reviewers



Ashish Soni is a qualified and experienced Statistical Consultant and Data Science Professional with a strong understanding of and enthusiasm for data collection, data consolidation, statistical analysis, data mining, Machine Learning, and Artificial Intelligence. Over his professional career span of 11+ years, he employed and acquired comprehensive knowledge of diverse statistical and machine learning models and tests, using different suites and platforms. Ashish holds a BTech degree in Chemical Engineering from Indian Institute of Technology, Bombay, India; a masters degree in Economics; and a postgraduate diploma in Applied Statistics. He has worked across different industries such as healthcare; banking, financial services, and insurance; sports; and human resources; and he is currently providing his services as Delivery Head, Analytics and Data Science at the healthcare analytics firm Equantx Pharma Analytics Solutions, Delhi, India.

ABOUT THE TECHNICAL REVIEWERS



Girisha Garg is a Post Doctorate from Georgia State University and holds a doctorate (Instrumentation & Control Engineering), from the Netaji Subash Institute of Technology, New Delhi, India. With over 10+ years of industry and research experience and over 10+ research publications and book chapters, she brings a wealth of experience in Machine Learning, Deep Learning, Feature Engineering and Data Science. She has been involved in teaching since 2007 and has also served as a

corporate trainer for AIML (Artificial Intelligence and Machine Learning) and faculty for various undergraduate engineering institutes. Her portfolio of companies include Georgia State University, Hemisphere Technocrats - Mumbai, Move Forward Technologies - Delhi, Babu Banarasi Das Institute of Technology - Ghaziabad, Netaji Subhas Institute of Technology - Delhi, National Institute of Technology - Delhi, I.I.I.T-D Delhi, Y.M.C.A Institute of Engineering, Faridabad.

Acknowledgments

This book would not have been possible without the patience, tolerance, and encouragement of many people. First and foremost, eternal gratitude to my parents and grandparents, who migrated to foreign lands carrying nothing but hope, endured unimaginable hardships, and gave us the world.

A huge thank you to the Diabetes Digital Media team—particularly Amar, Harkrishan, and Dom—with whom many discussions on the future of artificial intelligence (AI) were had. Thank you to Krystal for creating beautiful imagery to complement the book.

Thank you Ashish Soni and Girisha Garg, for the meticulous detail and technical rigor you brought to this book and the fantastic team at Apress—Divya, Celestin—for your encouragement, support, and motivation.

Finally, thank you to time: what a funny old thing. Dedicated to the infinite majesty of Kirpa, Ananta, and Charlotte.

Introduction

The world is changing. There are more phones than people in the world, and it is increasingly connected. People use virtual assistants, self-driving cars, find partners through digital apps, and search the Web for any symptom of ill health. Each digital event leaves a digital exhaust that is datafying life as we know it. The success of many of the world's most loved services, from Google to Uber, Alexa to Netflix, is grounded in big data and optimization.

Although medicine has been receptive to the benefits of big data and AI, it has been slow to adopt the rapidly evolving technology, particularly when compared to sectors such as finance, entertainment, and transport: that is, until now.

Recent digital disruption has catalyzed healthcare's adoption of big data and AI. Data of all sorts, shapes, and sizes are used to train AI technologies that facilitate machines to learn, adapt, and improve on their learning. Academic institutions and start-ups alike are developing rapid prototype technologies with increasingly robust health and engagement claims. The blending of technology and medicine has been expedited by smartphones and the Internet of Things, which is facilitating a wealth of innovation that continues to improve lives. With the arrival of health technology, people can monitor their health without the assistance of a healthcare professional; healthcare is now mobile and no longer in the waiting room.

At the same time, the world's population is living longer, and unhealthier, than ever: and in a financial crisis. Healthcare services are turning to value-based and incentivized care as non-communicable diseases such as type 2 diabetes and obesity become global pandemics.

INTRODUCTION

Data is proving an invaluable tool in improving health; information is empowering. Data science, the science of big data, its analysis, and intelligent programming layers has now become a pillar to achieve traction and success in healthcare. Digital health can democratize and personalize healthcare—and data is the golden key. This comes at the same time as a growing appetite to measure and quantify more aspects of human life.

Data insight and real-world evidence are facilitating rapid technology innovation that regulators are struggling to keep up with. The consequences of digital health democratization have not only a health impact but also ethical implications. Big data and machine learning enable stakeholders to uncover hidden connections and patterns, including predictions of the future. The effects of understanding this data have moral and legal consequences that require appropriate governance to mitigate risk and harm.

Healthcare providers, individuals, and organizations all house a fountain of data that can be used for machine learning. Many people have a fair idea of what they would like to learn from data but are unaware as to how much data is required and what can be achieved before more technical aspects of uncovering hidden patterns, trends, and biases are found.

This book takes a practical, hands-on approach to big data, AI, and machine learning and the ethical implications of such tools. We cover the theory and practical applications of AI in healthcare—from where and how to start to applying machine learning techniques and evaluating performance. The book concludes with a series of case studies from leading health organizations who utilize AI and big data in novel and innovative ways.

CHAPTER 1

What Is Artificial Intelligence?

“Knowledge on its own is nothing, but the application of useful knowledge? That’s powerful.”

—Osho

Artificial intelligence (AI) is considered, once again, to be one of the most exciting advances of our time. Virtual assistants can determine our music tastes with remarkable accuracy, cars are now able to drive themselves, and mobile apps can reverse diseases once considered to be chronic and progressive.

Many people are surprised to learn that AI is nothing new. AI technologies have existed for decades. It is, in fact, going through a resurgence—and it is being driven by availability of data and cheaper computing.

A Multifaceted Discipline

AI is a subset of computer science that has origins in mathematics, logic, philosophy, psychology, cognitive science, and biology, among others (Figure 1-1).

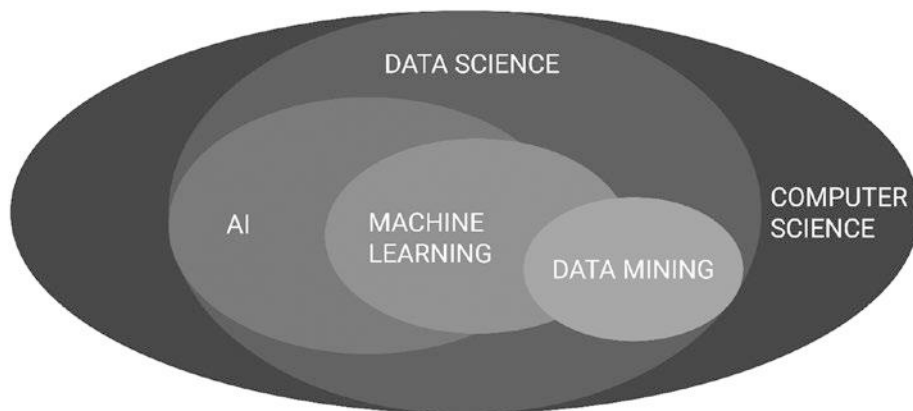


Figure 1-1. *AI, machine learning, and their place in computer science*

The earliest research into AI was inspired by a constellation of thought that began in the late 1930s and culminated in 1950 when British pioneer Alan Turing published *Computing Machinery and Intelligence* in which he asked, can machines think? The Turing Test proposed a test of a machine's ability to demonstrate “artificial” intelligence, evaluating whether the behavior of a machine is indistinguishable from that of a human. Turing proposed that a computer could be considered to be able to think if a human evaluator could have a natural language conversation with both a computer and a human and not distinguish between either (i.e., an agent or system that is successfully mimicking human behavior).

The term AI was first coined in 1956 by Professor John McCarthy of Dartmouth College. Professor McCarthy proposed a summer research project based on the idea that “every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it”.[1]

The truth is that AI, at its core, is merely programming. As depicted in Figure 1-1, AI can be understood as an abstraction of computer science. The surge in its popularity, and so too its ability, has much to do with the

explosion of data through mobile devices, smartwatches, wearables, and the ability to access computer power cheaper than ever before. It was estimated by IBM in 2011 that 90% of global data had been created in the preceding 2 years.[2]

It is estimated that there will be 150 billion networked measuring sensors in the next decade—which is 20 times the global population. This exponential data generated is enabling everything to become smart. From smartphones to smart washing cars, smart homes, cities, and communities await.

With this data comes a plethora of learning opportunities and hence, the focus has now shifted to learning from available data and the development of intelligent systems. The more data a system is given, the more it is capable of learning, which allows it to become more accurate.

The use and application of AI and machine learning in enterprise are still relatively new, and even more so in health. The Gartner “Hype Cycle for Emerging Technologies” in 2017 placed machine learning in the peak of inflated expectations, with 5 to 10 years before plateau.

As a result, the applications of machine learning within the healthcare setting are fresh, exciting, and innovative. With more mobile devices than people today, the future of health is wrought with data from the patient, environment, and physician. As a result, the opportunity for optimizing health with AI and machine learning is ripening.

The realization of AI and machine learning in health could not be more welcome in the current ecosystem, as healthcare costs are increasing globally, and governmental and private bill payers are placing increasing pressures on services to become more cost-effective. Costs must typically be managed without negatively impacting patient access, patient care, and health outcomes.

But how can AI and machine learning be applied in an everyday healthcare setting? This book is intended for those who seek to understand what AI and machine learning are and how intelligent systems can be

developed, evaluated, and deployed within their health ecosystem. Real-life case studies in health intelligence are included, with examples of how AI and machine learning is improving patient health, population health, and facilitating significant cost savings and efficiencies.

By the end of the book, readers should be confident in explaining key aspects of AI and machine learning to stakeholders. Readers will be able to describe the machine learning approach and limitations, fundamental algorithms, the usefulness and requirements of data, the ethics and governance of learning, and how to evaluate the success of such systems.

Rather than focus on overwhelming statistics and algebra, theory and practical applications of AI and machine learning in healthcare are explored—with methods and tips on how to evaluate the efficacy, suitability, and success of AI and machine learning applications.

Examining Artificial Intelligence

At its heart, AI can be defined as the simulation of intelligent behavior in agents (computers) in a manner that we, as humans, would consider to be smart or human-like. The core concepts of AI include agents developing traits including knowledge, reasoning, problem-solving, perception, learning, planning, and the ability to manipulate and move.

In particular, AI could be considered to comprise the following:

- Getting a system to reason rationally. Techniques include automated reasoning, proof planning, constraint solving, and case-based reasoning.
- Getting a program to learn, discover and predict. Techniques include machine learning, data mining (search), and scientific knowledge discovery.

- Getting a program to play games. Techniques include minimax search and alpha-beta pruning.
- Getting a program to communicate with humans. Techniques include natural language processing (NLP).
- Getting a program to exhibit signs of life. Techniques include genetic algorithms.
- Enabling machines to navigate intelligently in the world
- This involves robotic techniques such as planning and vision.

There are many misconceptions of AI, primarily as it's still quite a young discipline. Indeed, there are also many views as to how it will develop. Interesting expert opinions include those of Kevin Warwick, who is of the opinion robots will take over the earth. Roger Penrose reasons that computers can never truly be intelligent. Meanwhile, Mark Jeffery goes as far as to suggest that computers will evolve to be human. Whether AI will take over the earth in the next generation is unlikely, but AI and its applications are here to stay.

In the past, the intelligence aspect of AI has been stunted due to limited datasets, representative samples of data, and the inability to both store and subsequently index and analyze considerable volumes of data. Today, data comes in real time, fuelled by exponential growth in mobile phone usage, digital devices, increasingly digitized systems, wearables, and the Internet of Things (IoT).

Not only is data now streaming in real time, but it also comes in at a rapid pace, from a variety of sources, and with the demand that it must be available for analysis, and fundamentally interpretable, to make better decisions.

There are four distinctive categories of AI.

Reactive Machines

This is the most basic AI. Reactive systems respond in a current scenario, relying on taught or recalled data to make decisions in their current state. Reactive machines perform the tasks they are designed for well, but they can do nothing else. This is because these systems are not able to use past experiences to affect future decisions. This does not mean reactive machines are useless. Deep Blue, the chess-playing IBM supercomputer, was a reactive machine, able to make predictions based on the chess board at that point in time. Deep Blue beat world champion chess player Garry Kasparov in 1996. A little-known fact is that Kasparov won three of the remaining five games and defeated Deep Blue by four games to two.

Limited Memory—Systems That Think and Act Rationally

AI that works off the principle of limited memory and uses both pre-programmed knowledge and subsequent observations carried out over time. During observations, the system looks at items within its environment and detects how they change, then makes necessary adjustments. This technology is used in autonomous cars. Ubiquitous Internet access and IoT is providing an infinite source of knowledge for limited memory systems.

Theory of Mind—Systems That Think Like Humans

Theory of mind AI represents systems that interpret their worlds and the actors, or people, in them. This kind of AI requires an understanding that the people and things within an environment can also alter their feelings and behaviors. Although such AI is presently limited, it could be used in caregiving roles such as assisting elderly or disabled people with everyday tasks.

As such, a robot that is working with a theory of mind AI would be able to gauge things within their worlds and recognize that the people within the environments have their own minds, unique emotions, learned experiences, and so on. Theory of mind AI can attempt to understand people's intentions and predict how they may behave.

Self-Aware AI—Systems That Are Humans

This most advanced type of AI involves machines that have consciousness and recognize the world beyond humans. This AI does not exist yet, but software has been demonstrated with desires for certain things and recognition of its own internal feelings. For instance, in 2015, researchers at the Rensselaer Polytechnic Institute gave an updated version of the wise men puzzle, an induction self-awareness test, to three robots—and one passed. The test requires AI to listen and understand unstructured text as well as being able to recognize its own voice and its distinction from other robots.

Technology is now agile enough to access huge datasets in real time, and learn on the go. Ultimately, AI is only as good as the data that's used to create it—and with robust, high-volume data, we can be more confident about our decisions.

Healthcare has been slow to adopt the benefits of big data and AI, especially when compared to transport, energy, and finance. Although there are many reasons for this, the rigidity of the medical health sector has been duly grounded in the fact that peoples' lives are at risk. Medical services are more of a necessity than a consumer choice; so historically, the medical industry has had little to no threat that usually drives other industries to seek innovation.

That has expedited a gap in what healthcare institutes can provide and what patients want—which subsequently has led to variances in care, in health outcomes, and a globally recognized medication-first approach to disease.

The explosion of data has propelled AI through enabling a data-led approach to intelligence. The last 5 years has been particularly disruptive in healthcare, with applications of data-led AI helping intelligent systems not only to predict, diagnose, and manage disease but to actively reverse and prevent it—and the realization of digital therapeutics.

Recent advances in image recognition and classification are beginning to filter into industry too, with deep neural networks achieving remarkable success in visual recognition tasks, often matching or exceeding human performance (Figure 1-2).

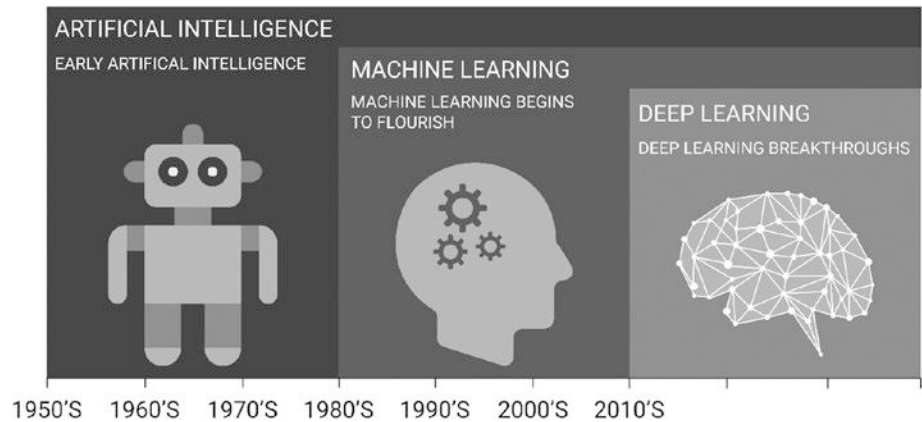


Figure 1-2. *AI and its development*

What Is Machine Learning?

Machine learning is a term credited to Arthur Samuel of IBM, who in 1959 proposed that it may be possible to teach computers to learn everything they need to know about the world and how to carry out tasks for themselves. Machine learning can be understood as an application of AI.

Machine learning was born from pattern recognition and the theory that computers can learn without being programmed to perform specific tasks. This includes techniques such as Bayesian methods; neural

networks; inductive logic programming; explanation-based, natural language processing; decision tree; and reinforcement learning.

Systems that have hard-coded knowledge bases will typically experience difficulties in new environments. Certain difficulties can be overcome by a system that can acquire its own knowledge. This capability is known as machine learning. This requires knowledge acquisition, inference, updating and refining the knowledge base, acquisition of heuristics, and so forth.

Machine learning is covered in greater depth in Chapter 3.

What Is Data Science?

All AI tasks will use some form of data. Data science is a growing discipline that encompasses anything related to data cleansing, extraction, preparation, and analysis. Data science is a general term for the range of techniques used when trying to extract insights (i.e., trying to understand) and information from data.

The term data science was phrased by William Cleveland in 2001 to describe an academic discipline bringing statistics and computer science closer together.

Teams working on AI projects will undoubtedly be working with data, whether little or big in volume. In the case of big data, real-time data usually demands real-time analytics. In most business cases, a data scientist or data engineer will be required to perform many technical roles related to the data including finding, interpreting, and managing data; ensuring consistency of data; building mathematical models using the data; and presenting and communicating data insights/findings to stakeholders. Although you cannot do big data without data science, you can perform data science without big data—all that is required is data.

Because data exploration requires statistical analysis, many academics see no distinction between data science and statistics. A data science team (even if it is a team of one) is fundamental to deploying a successful project.

The data science team typically performs two key roles. First, beginning with a problem or question, it is seeking to solve it with data; and second, it is using the data to extract insight and intelligence through analytics (Figure 1-3).

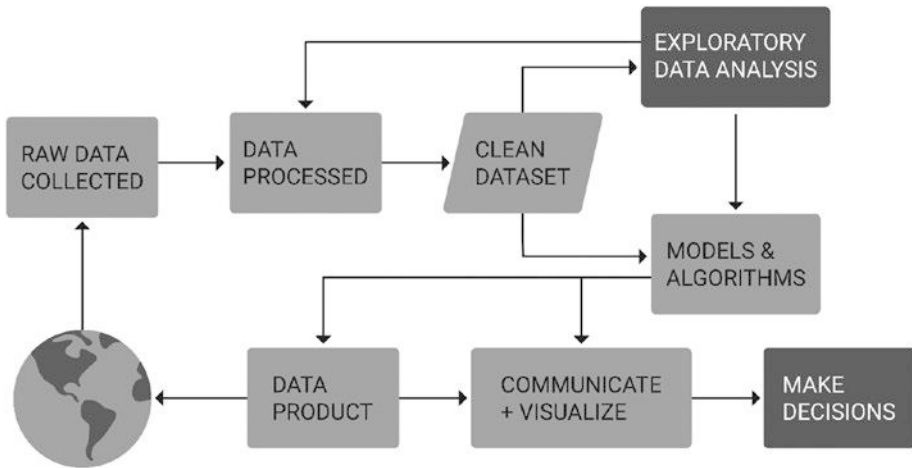


Figure 1-3. *Data science process*

Learning from Real-Time, Big Data

AI had previously been stifled by the technology and data at its disposal. Before the explosion of smartphones and cheaper computing, datasets remained limited in respect to their size, validity, and representative rather than real-time nature.

Today, we have real-time, immediately accessible data and tools that enable rapid analysis. Datafication of modern-day life is fuelling machine learning's maturity and is facilitating the transition to an evidence-based, data-driven approach. Datafication refers to the modern-day trend of digitalizing (or datafying) every aspect of life). Technology is now agile enough to access these huge datasets to rapidly evolve machine learning applications.

Both patients and healthcare professionals generate a tremendous amount of data. Phones collect metrics such as blood pressure, geographical location; steps walked, nutritional diaries; and other unstructured data such as conversations, reactions, and images.

It's not just digital or clinical data either. Data extraction techniques can be applied to paper documentation, or images scanned, to process the documents for synthesis and recall.

Healthcare professionals collect health biomarkers alongside other structured metrics. Regardless of the source of data, to learn, data must be turned into information. For example, an input of a blood glucose reading from a person with diabetes into an app has far more relevance when blood glucose level targets are known by the system so that it can be understood whether the input met the recommended target range or not.

In the twenty-first century, almost every action leaves some form of transactional footprint. The Apple iPhone 6 is over 32,000 times faster than the Apollo-era NASA computers that took man to the moon. Not only are computers smaller than ever, but they are also more powerful than ever. With organizations capitalizing on sources of big data, there has been a shift toward embedding learnings from data into every aspect of the user experience—from buying a product or service to the user experience within an app.

The value of data is understood when it is taken in its raw form and converted into knowledge that changes practice. This value is driven by project and context. For example, the value may be as the result of faster identification of shortfalls in adherence, compliance, and evidence-based care. It may be better sharing of data and insights within a hospital or organization or more customized relationships with patients to drive adherence, compliance, and boost self-care. Equally, it may be to avoid more costly treatments or costly mistakes.

Applications of AI in Healthcare

It is unlikely artificially intelligent agents will ever completely replace doctors and nurses, but machine learning and AI are transforming the healthcare industry and improving outcomes. Machine learning is improving diagnostics, predicting outcomes, and beginning to scratch the surface of personalized care (Figure 1-4).

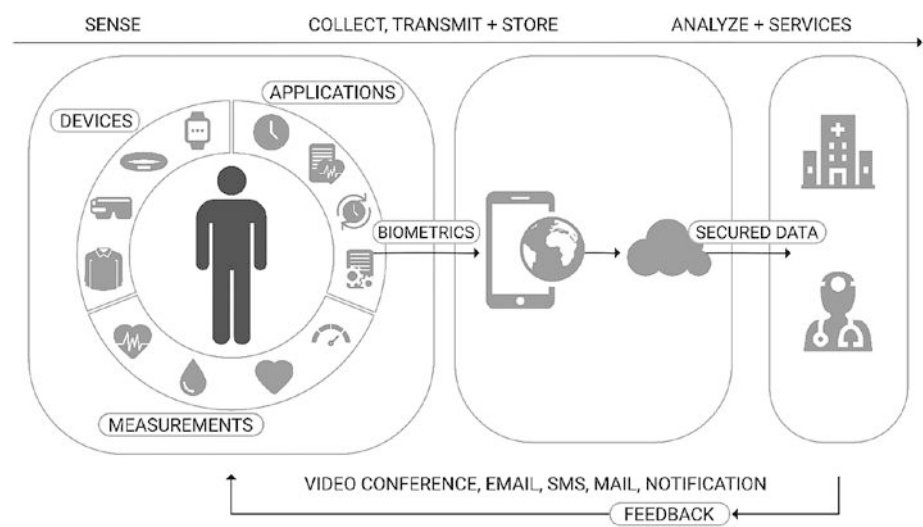


Figure 1-4. A data-driven, patient-healthcare professional relationship

Imagine a situation where you walk in to see your doctor with pain around your heart. After listening to your symptoms, she inputs them into her computer, which pulls up the latest evidence base she should consult to efficiently diagnose and treat your problem. You have an MRI scan, and an intelligent computer system assists the radiologist in detecting any concerns that would be too small for your radiologist’s human eye to see.

Your watch may have even been continuously collecting your blood pressure and pulse while a continuous blood glucose monitor had a

real-time profile of your blood glucose readings. Finally, your medical records and family's medical history is assessed by a computer system that suggests treatment pathways precisely identified to you. Data privacy and governance aside, the implications of what we can learn from combining various pools of data are exciting.

By the end of this book, readers will be able to lead the development of their own machine learning project.

Prediction

Technologies already exist that monitor data to predict disease outbreaks. This is often done using real-time data sources such as social media as well as historical information from the Web and other sources. Malaria outbreaks have been predicted with artificial neural networks, analyzing data including rainfall, temperature, number of cases, and various other data points.

Diagnosis

Many digital technologies offer an alternative to non-emergency health systems. Considering the future, combining the genome with machine learning algorithms provides the opportunity to learn about the risk of disease, improve pharmacogenetics, and provide better treatment pathways for patients.[3]

Personalized Treatment and Behavior Modification

A digital therapy from Diabetes Digital Media, the Low Carb Program, assists people with type 2 diabetes and prediabetes to reverse (i.e., place into remission) their condition. The app provides personalized education

and integrated health tracking, learning from the user's and wider community's progress. At the end of a year, most members who complete the program reduce medication dependency, saving over \$1,015 per patient per year in medication "deprescription." [4]

Drug Discovery

The use of machine learning in preliminary drug discovery has the potential for various uses, from initial screening of drug compounds to predicted success rate based on biological factors. This includes R&D discovery technologies like next-generation sequencing.

Drugs do not necessarily need to be pharmacological in their appearance. The use of digital solutions and aggregation of real-world patient data are providing solutions to conditions once considered to be chronic and progressive. The Low Carb Program app for example, used by over 300,000 people with type 2 diabetes places the condition into remission for 26% of the patients who complete the program at 1-year. [5]

Follow-Up Care

Hospital readmittance is a huge concern in healthcare. Doctors, as well as governments, are struggling to keep patients healthy, particularly when returning home following hospital treatment. Organizations such as NextIT have developed digital health coaches, similar to a virtual customer service representative on an e-commerce site. The assistant prompts questions about the patient's medications and reminds them to take medicine, queries them about their condition symptoms, and conveys relevant information to the doctor.

Realizing the Potential of AI in Healthcare

For AI and machine learning to be fully embraced and integrated within healthcare systems, several key challenges must be addressed.

Understanding Gap

There is a huge disparity between stakeholder understanding and applications of AI and machine learning. Communication of ideas, methodologies, and evaluations are pivotal to the innovation required to progress AI and machine learning in healthcare. Encouraging the adoption of data-driven strategies Data, including the sharing and integration of data, is fundamental to shift healthcare toward realizing precision medicine.

Developing data science teams, focused on learning from data, is key to a successful healthcare strategy. The approach required is one of investment in data. Improving value for both the patient and the provider requires data and hence data science professionals.

Fragmented Data

There are many hurdles to be overcome. Data is currently fragmented and difficult to combine. Patients collect data on their phones, Fitbits, and watches, while physicians collect regular biomarker and demographic data. At no point in the patient experience is this data combined. Nor do infrastructures exist to parse and analyze this larger set of data in a meaningful and robust manner. In addition, electronic health records (EHRs), which at present are still messy and fragmented across databases, require digitizing in a mechanism that is available to patients and providers at their convenience.

Appropriate Security

At the same time, organizations face challenges of security and meeting government regulation specifically with regards to the management of patient data and ensuring its accessibility at all times. What's more, many healthcare institutions are using legacy versions of software that can be more vulnerable to attack. The NHS (National Health Service) digital infrastructure was paralyzed the wrath of the ransomware WannaCry in 2017. The ransomware, which originated in America, scrambled data on computers and demanded payments of \$300 to \$600 to restore access.[6]

Hospitals and GP surgeries in England and Scotland were among over 16 health service organizations hit by the "ransomware" attack. The impact of the attack wasn't just the cost of the technological failure: it had a bearing on patients' lives. Doctors' surgeries and hospitals in some parts of England had to cancel appointments and refuse surgeries. In some areas, people were advised to seek medical care in emergencies only. The NHS was just one of many global organizations crippled through the use of hacking tools; and the ransomware claimed to have infected computers in 150 countries.

Data Governance

With this comes the topic of data governance. Medical data is personal and not easy to access. It is widely assumed that the general public would be reluctant to share their data because of privacy concerns. However, a Wellcome Foundation survey conducted in 2016 on the British public's attitude to commercial access to health data found that 17% of people would never consent to their anonymized data being shared with third parties.[7]

Adhering to multiple sets of regulation means disaster recovery and security is key, and network infrastructure plays a critical role in ensuring these requirements can be met. Healthcare organizations require

modernization of network infrastructure to ensure they are appropriately prepared to provide the best patient care possible. The case for this is highlighted by the fact that the bulk of NHS computers in 2018 use Internet Explorer 8 as their default Internet browser - which was first launched almost a decade ago.[8]

Bias

A significant problem with learning is bias. As AI becomes increasingly interwoven into our daily lives—integrated with our experiences at home, work, and on the road—it is imperative that we question how and why machines do what they do. Within machine learning, learning to learn creates its own inductive bias based on previous experience. Essentially, systems can become biased based on the data environments they are exposed to.

It's not until algorithms are used in live environments that people discover built-in biases, which are often amplified through real-world interactions.

This has expedited the growing need for more transparent algorithms to meet the stringent regulations on drug development and expectation. Transparency is not the only criteria; it is imperative to ensure decision-making is unbiased to fully trust its abilities. People are given confidence through the ability to see through the black box and understand the causal reasoning behind machine conclusions.

Software

Traditional AI systems had been developed in Prolog, Lisp, and ML. Most machine learning systems today are written in Python due to many of the mathematical underpinnings of machine learning that are available as libraries. However, algorithms that “learn” can be developed in most languages, including Perl, C++, Java, and C. This is discussed in more depth in Chapter 3.

Conclusion

The potential applications of machine learning in healthcare are vast and exciting. Intelligent systems can help us reverse disease, detect our risk of cancers, and suggest courses of medication based on real-time biomarkers.

With this also comes tremendous responsibility and questions of wider morality. We don't yet fully understand what can be learned from health data. As a result, the ethics of learning is a fundamental topic for consideration.

On the basis that an intelligent system can detect the risk of disease, should it tell the patient it is tracking the impending outcome? If an algorithm can detect your risk of pancreatic cancer—an often-terminal illness—based on your blood glucose and weight measurements, is it ethical to disclose this to the patient in question? Should such sensitive patient data be shared with healthcare teams - and what the unintended consequences of such actions?

The explosion in digitally connected devices and applications of machine learning are enabling the realization of these once-futuristic concepts, and conversation around these topics is key to progress.

And if a patient opts out of sharing data, is it then ethical to generalize based on known data to predict the same illness risk? And what if I can't get life insurance without this pancreatic cancer check? There are considerable privacy concerns associated with the use of a person's data and what should be private or not, and equally as to what data is useful. Invariably, the more data available, the more precise a decision can be made—but exactly how much is too much is another question. The driving factor is determining the value of data, which will be discussed further in Chapter 2.

The ethics of AI are currently without guidelines, regulations, or parameters on how to govern the enormous treasure chest of data and opportunity. Many assume that AI has an objectivity that puts it above

questions of morality, but that's not the case. AI algorithms are only as fair and unbiased as the learnings, which come from the environmental data. Just like social relationships, political and religious affiliations, and sexual orientation can be revealed from data, learning on health data is revealing new ethical dilemmas for consideration.

Data governance and disclosure of such data still requires policy, at the national and international level. In the future, driverless cars will be able to use tremendous amounts of data in real time to predict the likelihood of survival. Would it be ethical for the systems to choose who lives or dies or for a doctor to decide whom to treat based on the reading from two patients' Apple Watches?

In reality, this is just the beginning. As technologies develop, new and improved treatments and diagnoses will save more lives and cure more diseases. With the future of medicine grounded in data and analytics, it begs the question as to whether there will ever be enough data.

CHAPTER 2

Data

“The plural of anecdote is data.”

—Raymond Wolfinger

Data is everywhere. Global disruption and international initiatives are driving datafication. Datafication refers to the modern-day trend of digitalizing (or datafying) every aspect of life. This data creation is enabling the transformation of data into new and potentially valuable forms. Entire municipalities are being incentivized to become smarter. In the not too distant future, our towns and cities will collect thousands of variables in real time to optimize, maintain, and enhance the quality of life for entire populations. One would reasonably expect that as well as managing traffic, traffic lights may also collect other data such as air quality, visibility, and speed of traffic. As a result of big data from connected devices, embedded sensors, and the IoT, there is a global need for the analysis, interpretation, and visualization of data.

What Is Data?

Data itself can take many forms—character, text, words, numbers, pictures, sound, or video. Each piece of data falls into two main types: structured and unstructured.

At its heart, data is a set of values of qualitative or quantitative variables. To become information, data requires interpretation. Information is organized or classified data, which has some meaningful value (or values) for the receiver. Information is the processed data on which decisions and actions should be based.

Healthcare is undergoing a data revolution, thanks to two huge shifts: the need to tame growing cost pressures, which is generating new incentives and reimbursement structures, and digital health, which is democratizing healthcare by empowering consumers through their digital devices.

Clinical trends are changing. The advent of social media and immediate access to information through the Internet and health communities mean that patients are clued up about their health. Patients are generating data constantly and want to own and transfer their data between services. Concurrently, medicine is striving for personalized care, led through an evidence-based approach to decision-making. Patients and healthcare professionals alike are vocal in their appetite for all available clinical data to make more effective treatment decisions based on current evidence. A Diabetes.co.uk study in 2018 demonstrated that one in three patients would like to share their data with their healthcare professional, but only one in five patients can do so.[9]

Aggregating individual datasets into more significant populations also provides more robust evidence because subtleties in subpopulations may be infrequent in that they are not readily apparent in small samples. For instance, aggregating patients with type 1 and type 2 diabetes may present health conditions such as retinopathy with more confidence than separate datasets.

The incorporation of data in modern healthcare provides an opportunity for significant improvements within a plethora of areas. As medicine evolves to become evidence based and personalized, data use can be used to improve

- Patient and population wellness
- Patient education and engagement
- Prediction of disease and care risks
- Medication adherence
- Disease management
- Disease reversal/remission
- Individualization and personalization of treatment and care
- Financial, transactional, and environmental forecasting, planning, and accuracy

Data facilitates information. Information leads to insights, and insights lead to the making of better decisions.

Types of Data

Structured data typically refers to something stored in a database—in a structure that follows a model or schema. Almost every organization will be familiar with this form of data and may already be using it effectively. Most organizations will store at least some form of data in Excel spreadsheets, for example. Within a clinical setting, EHRs may also take a similar, structured form. Readings from embedded sensors, smartphones, smartwatches, and IoT devices are typically forms of structured data—whether that be the provision of blood glucose readings, steps walked, calories burned, heart rate, or blood pressure.

Structured data is similar to a machine language. Highly organized in its format, structured data facilitates simple, straightforward search and information retrieval operations. Structured data would typically be stored in a relational database for this purpose.

Unstructured data refers to everything else. Unstructured data does not have a predefined model or schema. Data that is unstructured has no identifiable structure within it, and this presents problems for querying and information retrieval. E-mails, text messages, Facebook posts, Twitter tweets, and other social media posts are good examples of unstructured data.

The Gartner report has indicated that data volume is set to grow 800% over the next 5 years, and 80% of this data will be in the form of unstructured data.[10]

The problem that unstructured data presents is one of volume. The lack of structure makes compilation and interpretation a resource-intensive task in regard to time and energy. Any organization would benefit from a mechanism of unstructured data analysis to reduce the resource costs unstructured data add to any organization.

Structured data is traditionally easier to analyze than unstructured data due to unstructured data's raw and unorganized form. However, analytics of unstructured, informal datasets is improving, with the use of data science and machine learning methods such as Natural Language Processing (NLP) assisting in the understanding and classification of sentiment.

It may not always be possible to transform unstructured data into a structured model. For instance, the transmission of information through an e-mail or notification holds data such as the time sent, subject, and sender as uniform fields. However, the contents of the message would not be easily dissected and categorized.

Semi-structured data fits the space between structured and unstructured data. Semi-structured data does not necessarily conform with the formal structures of data schemas associated with relational databases or data tables. However, semi-structured data may contain tags or markers to separate semantic elements and enforce grouping of records and fields within the data. Languages such as JSON (JavaScript Object Notation) and XML (Extensible Markup Language) are all forms of semi-structured data.

Data is the fuel required to learn in an intelligent system, the most critical component for successful AI. It is categorized into five sources:

- Web and social media data—clicks, history, health forums
- Machine-to-machine data—sensors, wearables
- Big transaction data—health claim data, billing data
- Biometric data—fingerprints, genetics, biomarkers
- Human-generated data—e-mail, paper documents, electronic medical records

When thinking about typical Excel spreadsheets, and in the domain of machine learning, the following definitions are also useful:

- Instance: A single row of data or observation.
- Feature: A single column of data. It is a component of the observation.
- Data type: This refers to the kind of data represented by the feature (e.g., Boolean, string, number)
- Dataset: A collection of instances used to train and test machine learning models.
- Training dataset: Dataset used to train the machine learning model
- Testing dataset: Dataset used to determine accuracy/performance of the machine learning model.

Big Data

The term big data is given to the collection of voluminous, traditional, and digital data that are sources for discovery and analysis. Big data is a popular term that defines datasets that are too big to be stored and processed in a conventional relational database system. In this way, the term big data is vague—while size is undoubtedly a part of big data, scale alone doesn't tell the whole story of what makes big data truly big.

Through the analytics of big data, we can uncover hidden patterns, unknown correlations, trends, preferences, and other information that can help stakeholders make better and more informed decisions. Machine learning provides a toolbox of techniques that can be applied to datasets for this very purpose.

Big data was first described by Laney in 2001 as having the following characteristics, also known as the 3 Vs (Figure 2-1).[11]

- **Volume**—the quantity of generated and stored data.
Big data is typically high in volume. The sheer size of data brings with it its intricacies in the form of storage, indexing and, retrieval.
- **Variety**—big data is varied in the types and nature of data, requiring efficient storage and analysis as well as systems for processing such data.
- **Velocity**—big data is received at a speed that brings its demands and challenges.

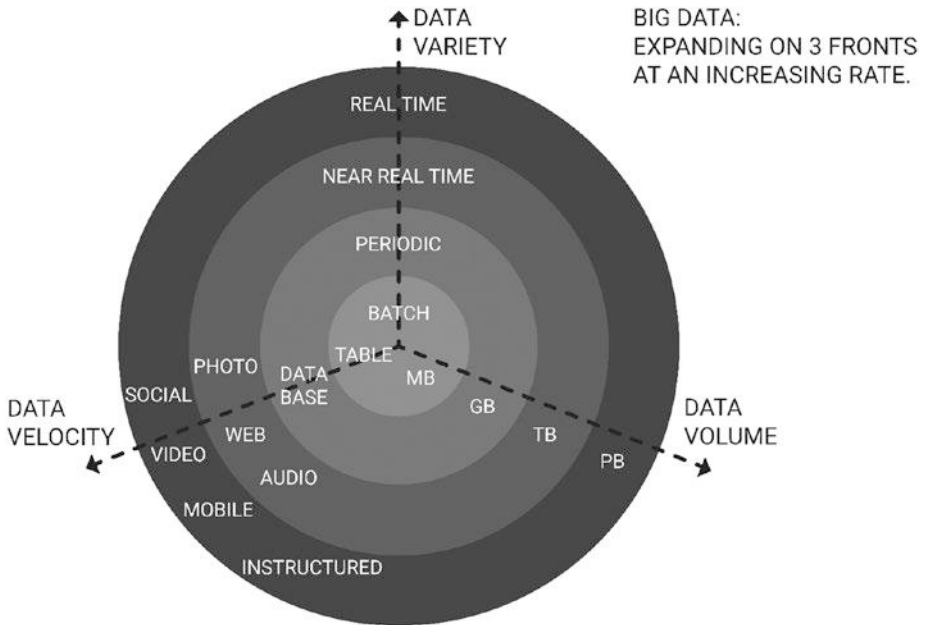


Figure 2-1. *The 3 Vs of data*

Data science has been discussing the 3 Vs of big data for some time. However, there are two other aspects of data to consider that are perhaps more important than the 3 Vs of big data: the concepts of data veracity and data value. An article by Mark van Rijmenam proposed 4 more Vs to further understand the incredibly complex nature of big data.[12]

Further still, there are a total of 10 Vs (Figure 2-2).[13] In reality, differing subsets of Vs are important for organizations to keep in mind when developing a data strategy. A subset of the many Vs now proposed have been included here.

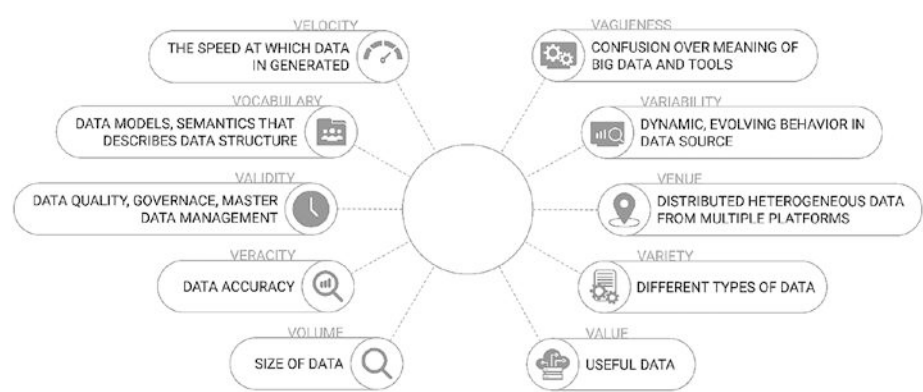


Figure 2-2. *The 10 Vs of data*

In 2016, a Gartner survey of 199 members demonstrated that investment in big data was increasing; but the investigation showed signs of slowing growth, with fewer companies having a future intent to invest in big data initiatives. Only 15% of businesses reported deploying their big data project, indicating that the adoption curve is still maturing and that talent and expertise is in high demand.[14]

Patients and physicians are driving the adoption of big data. Patients drive data through clinical data, wearable devices, mobile health apps, and telehealth remote healthcare services. Physicians meanwhile leave an exhaust of clinical data (such as written notes, imaging, and insurance data), patient records, and machine-generated data.

Volume

Laney defined a key characteristic of big data in 2001—the proposition that there is a lot of it. Key inertia in big data adoption has historically been the associated demand for storage and its respective costs. The movement of data centers from being locally housed to existing in the cloud has quashed this concern. This has significantly reduced storage costs as well as providing flexibility in data storage, collaboration, and

disaster recovery. The cloud, or cloud computing, refers to distributed computing delivered through the Internet.

Moore's Law has not just affected storage capacity; it has also benefitted cost. Moore's Law is a statement made by Gordon Moore, the Intel cofounder in 1965. Moore states that the number of transistors that are able to fit onto an integrated circuit doubles approximately every 18 months. In 1981, the price of a gigabyte of storage was USD 300,000. In 2004, this was \$1.00; and it was \$0.10 in 2010. Today, 1GB can be rented in the cloud for \$0.023 per month, with the first year of storage free.[15]

Healthcare, whether that be public health services, governments, pharmaceutical organizations, or global billpayers, is no stranger to big data. Data is continuously collected both digitally and nondigitally and spans a variety of patient, behavioral, epidemiological, environmental, and medical information. Data is consumed, managed, and generated for use in patient care, transactional records, research, compliance, processes, and regulatory requirements.

As of 2018, a patient's EHR is not readily available digitally in a secure and transparent means. Many organizations are striving to place this information on the blockchain, which revolutionizes access to health data through creating decentralized network systems.

Patients and healthcare organizations are adopting smartphones, wearables such as smartwatches and fitness trackers, home sensors, intelligent personal assistants, and social networks within their environments. This data is plentiful, accessible, and determined to be reliable. However, it is currently not used by most healthcare providers, highlighting the following:

- Patient data generation happens outside of the healthcare system.
- Patient-generated data needs to be integrated and explored within a more extensive patient record or health timeline.

- Patient data extends beyond medical markers to more holistic markers including behavioral and environmental health.
- There is an enormous interest for healthcare organizations still to be realized from opportunity and cost benefits.

A recent initiative in the United Kingdom saw the prevalence of type 1 diabetes in children geographically mapped through the use of social media hashtag #facesofdiabetes and #letstalk. Although this data is strictly not clinical, it highlights the power of user data and how it can be in real time for interpretation. The sentiment of users was also profiled, with appropriate support provided to users through the use of machine learning tools to classify post types.[16]

Datafication is well underway in the finance, energy, and retail sectors. With the prevalence of the health wearables and mobile healthcare, every individual and organization has the potential to tap into this growing data volume.

Coping with Data Volume

For big data to work for your organization, it is vital that storage is efficient and cost-effective. Choosing the most appropriate storage solution is dependent on several factors:

- Approach. Data storage costs increase alongside storage requirements. Hence, determining whether specific data is to be collected or whether a catchall approach to data is required depends on the project and regulatory requirements. Low-cost cloud storage providers make catchall approaches feasible. Beginners to data science and machine learning may wish to determine a set of key data metrics for evaluation for the sake of simplicity.

- **Types of data.** Consider whether structured or unstructured data is to be stored—and the source: for example, audio or video content, text, images, and so forth. Video, audio, or image storage will require more resources compared to text.
- **Deployment.** Determine how the solution will be deployed, which could include locally, internally, or cloud-based.
- **Access.** Determine how the solution will be accessed: through an app, web interface, intranet, and so forth.
- **Operations.** This includes data structure, architecture, archiving, recovery, event logging, and legal requirements.
- **Future use.** This involves planning for further sources of data, extensions to the system, and potential future use cases.

The distributed nature of datasets has meant that inertia to the adoption of big data has previously referred to the fragmented nature of such systems.

Variety

The second V of data is variety. This refers not only to variations in the types of data but also sources and use cases. Twenty years ago, we used to store data in the form of spreadsheets and databases. Today, data may be in the form of photographs, sensor data, tweets, encrypted files, and so on. This variety of unstructured data creates problems for storage, mining, and analyzing data. This is also an area that machine learning can significantly assist humans in.

Data is available in the form of structured data (in the form of clinical records and results), unstructured data (e.g., in the form of communication and interaction data), and also semi-structured (e.g., an X-ray with written annotations). The ever-increasing digitalization of services and the broad adoption of wearables and cyber-physical devices is continuously enabling new and exciting sources of data for discovery and analysis. This enables the creation of new risk models powered by sensor data, clinical records, communication and engagement, demographics, and billing records—which can make accurate predictions and save precious time and resources.

The Internet of Things

The Internet of Things refers to the rapidly increasing variety of smart, interconnected devices and sensors and the volumes of data they generate (Figure 2-3). This data is transferred between devices, services, and ultimately, people. You may also hear similar terms such as the Industrial Internet of Things and the Healthcare Internet of Things, which refer to their respective vertical, or industry, within the Internet of Things.

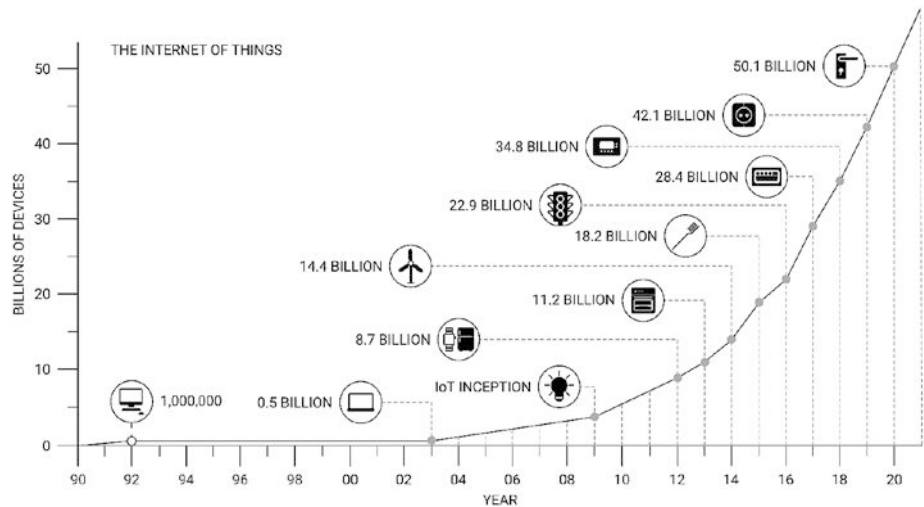


Figure 2-3. *Evolution of the Internet of Things*

Today, a variety of devices monitor every sort of patient behavior—from glucose monitors to fetal monitors to electrocardiograms to blood pressure. Many of these measurements require a follow-up visit with a healthcare professional. Smarter monitoring devices communicating with other devices, services, and systems could greatly refine this process, possibly lessening the need for direct clinical intervention and perhaps replacing it with a phone call from a nurse. For example, current innovations in smart devices include medication dispensers that can detect whether the medication has been taken—with data transmission occurring via Bluetooth. If it is determined that the user has not taken their medication, they are contacted via telephone to discuss and encourage adherence to their medication regime. There are a wealth of opportunities enabled by the Internet of Things to not only improve patient care but to reduce healthcare costs at the same time.

Variety of big data includes traditional data sources as well as newer sources of both structured and unstructured data. As the variety of data grows, so too does the potential of what algorithms and machine learning tools deployed in healthcare can achieve.

As well as a variety of sources, variety in big data can come in the form of the following:

- Data types—text, numbers, audio, video, images, and so forth.
 - Function—use cases and user requirements.
 - Value of data—is the data fit for a purpose? This would be a focus on quality in the context of the data's use.
- More data does not necessarily mean better data.

Just as there are a variety of data sources, value, types, and use cases, there is also a variety of applications for data. These include access points (i.e., web, mobile, SaaS [software as a service], API [application programming interface]) and users (typically either humans or machines).

Data variety is increasing exponentially. As new healthcare IoT devices develop, it will become imperative to identify useful data features from noise.

Legacy Data

Legacy data, whether computer-based or non-computer-based, can also be used in your projects. An increasing number of services enable digitalization for natural language processing or classification purposes. By using legacy data, stakeholders can maximize the insights from underutilized data. Legacy, or traditional, data may also be referred to as little or small data, discussed later in this chapter.

Typically, legacy data is fragmented and incomplete, particularly for the criteria desired today. For example, many records from more than 10 years ago may not have an associated e-mail address or correct phone number.

Velocity

The third V of big data is velocity, referring to the speed at which data is created, stored, and prepared for analysis and visualization. The velocity of data imposes unique demands on underlying computer hardware infrastructures.

A benefit of cloud computing has been the ability to quickly store and process the volume and variety of big data that would typically overwhelm a traditional server. Cloud computing is the preferred route for big data projects due to flexibility regarding storage and cost. Cloud providers can store petabytes of data and scale up thousands of servers in real time to requirements. Perhaps more valuable is that computational power is also inexpensive and distributable.

After the Haiti earthquake in 2010, Twitter data was a quicker way of detecting and tracking the deadly cholera outbreak compared to traditional methods. A subsequent research study determined that social media platforms outperformed official methods of monitoring disease prevalence in both speed and accuracy of detecting the progress of cholera.[17]

In the big data era, data is created in real time or near real time. The ubiquity of cyber-physical devices, embedded sensors, and other devices means that data transfer can occur at the moment it is created, bar a MAC address and Internet access.

The speed at which data is created is unimaginable. It is widely understood that the data generated in the previous 2 years is greater than the data created from the start of time up until then.[18] The challenge organizations have is to cope with the enormous speed at which data is created and consumed in real time.

As the penetration of wearable devices and sensors continues, so too will their application within clinical healthcare (Figure 2-4). The key to maximum clinical value is in the integration of various and heterogeneous data sources.

Data Type	Example	Characteristics
Patient behavior and sentiment	Social media Smartphones Web forums	Most data is unstructured or semi-structured. Data is large and in real time Opportunities for learning, engagement, and understanding. Typically open source
Patient health data	Sensors Blood glucose meters Smartphones Fitness trackers Images	Owned by patient. Data is typically device reported. Storage requirements typically low. Image storage requirements vary depending on data required for storage. Sensor data is high velocity. Structured data enables ease of classification and pattern detection. Some items may require digitalization

(continued)

Data Type	Example	Characteristics
Pharmaceutical and R&D data	Clinical trial data Population data	Encompasses population and disease data owned by pharmaceutical companies, research, academia, and the government
Health data on the Web	Patient portals (Diabetes.co.uk) Digital interventions (e.g., Low Carb Program, Gro Health, Hypo Program)	Useful metadata. Typically device reported rather than clinical
Clinical data	EHR Patient registries Physician data Images (scans, etc.)	Structured in nature. Owned by providers
Claims, cost, administrative data	Claims data Expenses data	Structured in nature. Encompasses any datasets relevant for reimbursement
Transactional data	Health information exchanges	May require preprocessing to become useful data

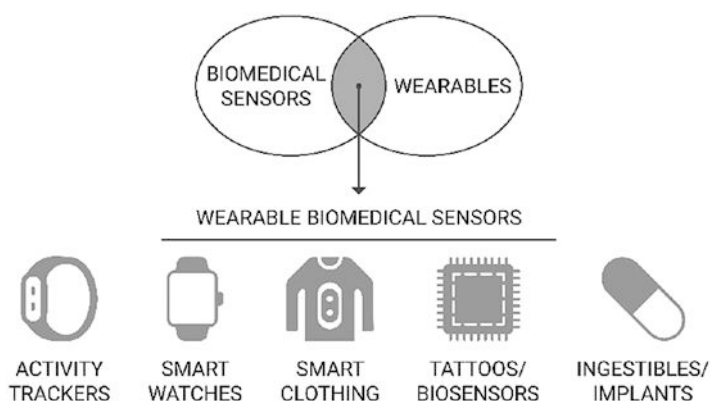


Figure 2-4. *Wearable biomedical sensors*

The traditional 3 Vs give an insight into the scale of data, and the rapid speeds at which these vast datasets grow and multiply. However, data variety only begins to scratch the surface of the depth and challenges of big data.

The power of big data was best demonstrated by monolith Google in 2009 when it was able to track the spread of influenza with only a one-day delay across the United States, faster than the Centers for Disease Control and Prevention (CDC), through the analysis of associated search terms.[19]

However, in 2013, Google’s Flu Trends got it wrong—which spurred questions on the concepts of value and validity of data.

In a clinical setting, data value and veracity would inherit an elevated position, as clinical decisions are made on data with lives at risk and go far beyond transactional implications.

Value

Value refers to the usefulness of data. Key to building trust in your data is ensuring its accuracy and usability. In healthcare, this would be evaluated through qualitative and quantitative analysis. Value can be circumscribed

through a variety of factors, primarily clinical impact on patient outcomes or behavior modification, patient and stakeholder engagement, impact on processes and workflows, and monetization (cost saving, cost benefits, etc.).

McKinsey stated that the potential annual value of big data to US Healthcare is \$300 billion. They also mention that big data has a potential annual value of €250 billion to Europe's public sector administration. Even more, in their well-regarded report from 2011, they state that the potential annual consumer surplus from using personal location data globally could be up to \$600 billion in 2020.[20]

Data's value comes from the analyses conducted on it. It is this analysis of data that enables it to be turned into information, with the aim of eventually turning it into knowledge. The value lies in how organizations use available data and become information centric and data-driven for purposes of decision-making. Value is not found, it is made. The key is to make the data meaningful to your users and your organization by managing it well.

Google did a perfect job in tracking the searches for flu symptoms in 2013. However, it had no way of knowing that its results weren't valid regarding predicting flu prevalence. Google Flu Trends, at least in this instance, wasn't providing valuable results for this use case. To add insult to injury, Google was never to know that real-world cases of flu failed to match the search distribution and frequency—so it could never do anything about it. Perhaps if this project was to be extended, analyzing social media posts may have been more valuable in determining flu prevalence. Social media and wider unstructured data sources provide a wealth of often unexplored information.

Identifying a population at risk of type 2 diabetes, predicting atrial fibrillation (AFIB), assisting patients with recovery from health events, or recommending the latest evidence-based medicine therapy to a patient are tasks with a clear and valuable return on investments.

Veracity

Worthless data, no matter how plentiful, is always going to be worthless. Veracity refers to data truthfulness and whether data is of optimal quality and suitability in the context of its use relating to the biases, noise, and abnormality in data.

Various factors can affect the veracity of data, including the following:

- **Data entry.** Was data entered correctly? Have there been any errors or events? Is there an audit trail of data entry?
- **Data management.** What is the integrity of the data moving through the system?
- **Integration quality.** Is data appropriately referenced, commented on, and unique?
- **Staleness.** Is data appropriate for use? Is it timely and relevant?
- **Usage.** Is the data available in a way that is actionable? Will the data be useful toward achieving business goals? Is it ethical to use this data?

The six Cs of trusted data can help assess whether the data you are using for your projects has veracity. It is essential to ensure your datasets are clean, complete, current, consistent, and compliant.

- **Clean data.** This is the result of good data quality procedures such as deduplication, standardization, verification, matching, and processing. Clean data enables robust decisions to be made on quality, non-tainted data. Clean data is the biggest challenge for user trust.

- **Complete data.** This is the result of consolidated data infrastructures, techniques, and processes that enable robust decision-making.
- **Current data.** Fresh data is normally considered more trustworthy over stale data. The question then becomes, at what point is data not current? As real-time, current data grows, deciphering the signal from noise can become a more challenging task.
- **Consistent data.** Data must be consistent (thus enabling machine readability). This is required for cross-compatibility of systems and applies to metadata also.
- **Compliant data.** Compliance regulations can come from various sources such as stakeholders, clients, legal legislation, or as the result of a new policy. Although many data infrastructures and standards exist, data science as a category is still coming up with first-time problems for which governance and legislation are required. Compliance can mean different things to internal and external stakeholders. Internally, there will be standards to ensure data is compliant with quality, security, and privacy procedures. All stakeholders need to trust that data has been accessed and distributed by internal and external regulations. Typically, organizations may have a governance board for data and information.

The sixth C is **collaborative data**, which refers to collaboration over data to ensure that data management and business management goals are aligned. The sixth C refers more to the approach to data than the data itself.

While most advocate complete veracity in data for clinical care, the gold standard of clinical care is harder to achieve in real-world scenarios with such a varied input of data. Enforcing data quality may be considered too difficult to achieve cost-effectively. In this case, training an intelligent system to learn to estimate for parameters not seen may be of use. Veracity can refer to ensuring sizeable training samples for rich model-building and validation, which empowers whole-population analytics.

Veracity can be applied to the analyses performed on the data and ensuring they are correct. Not only does data need to be trustworthy, but so too do the algorithms and systems interpreting it.

Data entry can be a risky point that goes unnoticed without good data science. This is typically a human-based problem. Even in small data settings, humans make mistakes. To be clinically relevant, data requires maximum veracity. As the veracity of data improves, machine learning on this data enables more veracious conclusions.

Validity

Similar to big data veracity, there is also an issue of data validity, referring to whether is data is correct and accurate for the intended use. In clinical applications, validity may be considered to be the prioritized V—to ensure that only useful and relevant data is used. Truthfulness or veracity of data is absolute, whereas validity is contextual.

Variability

Big data is variable. Variability defines data where the meaning is regularly changing. Variability is very relevant in performing sentiment analyses. For example, in a series of tweets, a single word can have a completely different meaning.

Variability is often confused with variety. To illustrate, a florist may sell five types of roses. That is variety. Now, if you go to the florist for two weeks in a row and buy the same white rose ever day, each day it will have a subtly different form and fragrance. That is variability.

To perform proper sentiment analyses, algorithms need to be able to understand the context of texts and be able to decipher the exact meaning of a word in that particular context. This is still very difficult, even with progress in natural language processing abilities.

Visualization

Visualization is the final of the 7 Vs, referring to the appropriate analyses and visualizations required on big data to make it readable, understandable, and actionable. Visualization may not sound complex; however, overcoming the challenge of visualizing complex datasets is crucial for stakeholder understanding and development. Quite often, it is data visualization that becomes the principal component in transferring knowledge learned from datasets to stakeholders.

Small Data

Little data (or small data) is a direct contrast to big data. Big data is distributed, varied, and comes in real time; whereas small data is data that is accessible, informative, and actionable as a result of the format and volume. Examples of small data include patient medical records, prescription data, biometric measurements, a scan, or even Internet search histories. In comparison to organizations and services such as Google and Amazon, the amount of data is far smaller in comparison.

Analytics has gone through a rapid evolution spurred by the adoption and demands of big data. Going back to traditional datasets and applying

more modern techniques such as machine learning should not be overlooked and is a great place for a data scientist to start. It's not the size of the data; it's what you do with it that counts.

Metadata

Metadata is data about data—which is descriptive data about each asset, or individual piece, of data. Metadata provides granular information about a single file supporting the facility to discover patterns and trends from metadata as well as the data it is supporting.

Metadata gives information about a file's origin, date, time, and format; and it may include notes or comments. It is key to investing resources, predominantly time, in strategic information management to make sure assets are correctly named, tagged, stored, and archived in a taxonomy consistent with other assets in the collection. This facilitates quicker dataset linkage and maintains consistent methodology for asset management to ensure files are easy to find, retrieve, and distribute.

Big data hype has overshadowed metadata. However, with big data comes big metadata, enabling organizations to generate knowledge and leverage value. For example, Google and Facebook use taxonomy languages such as Open Graph[21] that help create a more structured web, enabling more robust and descriptive information to be provided to users. This, in turn, provides human-friendly results to users, optimizing click through and conversion.

Metadata can be very useful in any data-led project. For instance, a machine learning algorithm could use the metadata belonging to a piece of music, rather than (or alongside) the actual music itself, to be able to suggest further relevant music. Features of the music such as genre, artist, song title, and year of release could be obtained from metadata for relevant results.

Healthcare Data—Little and Big Use Cases

Healthcare stakeholders understand they are surrounded by masses of data from patients, professionals, and transactions. It is key to know how to drive value and meet KPIs (key performance indicators). The following are a selection of exciting healthcare data use cases.

Predicting Waiting Times

In Paris, France, four of the hospitals that comprise the Assistance Publique-Hôpitaux de Paris (AP-HP) teamed up with Intel and used data from internal and external sources, including 10 years of hospital admissions records, to determine day- and hour-based predictions of the number of patients expected to enter their facility.[22]

Time series analysis techniques were used to predict admission rates at different times. This data was made available to all surgeries and clinics and demonstrates an immediate way data could be used to improve efficiency and empower stakeholders.

Most, if not all, clinics across the world have access to similar data, which demonstrates just how healthcare is only beginning to scratch the service of data's application.

Reducing Readmissions

The same approach to waiting times can be used to help manage hospital costs. Using data analytics, at-risk patient groups can be identified based on medical history, demographics, and behavioral data. This can be used to provide the necessary care to reduce readmission rates. At UT Southwestern hospital in the United States, EHRs analytics led to a drop in the readmission rate of cardiac patients from 26.2% to 21.2% through successful identification of at-risk patients.[23]

Predictive Analytics

The preceding examples could very well use static data (i.e., non-real-time small data) and be fairly accurate in predicting waiting times and readmission intervals. The same concept of analyzing data can be used at scale for prediction of disease and the democratization of care.

In the United States, Optum Labs has collected the EHR for over 30 million patients, creating a database for predictive analytics tools to improve the delivery of care. The intention is to enable doctors to make data-driven, informed decisions with proximity and therefore improve patients' treatment.[24]

The robustness that 30 million health records provide allows models to be trained and validated to find people who fit predictive risk trends for certain diseases such as hypertension, type 2 diabetes, heart disease, and metabolic syndrome.

By analyzing patient data including age, social, economic demographics, fitness, and other health biomarkers, providers can improve care at both an individual and population level through not only predicting risk but through the delivery of treatments for optimal patient outcomes.

Electronic Health Records

EHRs haven't quite come to fruition as yet. The idea is theoretically simple: that every patient has a digital health record consisting of their details, demographics, medical history, allergies, clinical results, and so forth. Records can be shared, with patient consent, via secure computer systems and are available for healthcare providers from both public and private sectors. Each record comprises one modifiable file, which means that doctors can implement changes over time with no danger of data replication or inconsistencies.

EHRs make perfect sense; however, complete implementation across a nation is proving a task. In the United States, up to 94% of hospitals use EHRs according to HITECH research. Europe is further behind. A European Commission directive has set the task of creating a centralized European health record system by 2020.[25]

In the United States, Kaiser Permanente has implemented a system that shares data across all their facilities and made it easier to use EHRs. A McKinsey report highlighted how the data sharing system achieved an estimated \$1 billion in savings as the result of reduced office visits and lab tests. The data sharing system improved outcomes in cardiovascular disease.[26]

The EHR is evolving into the blockchain, which seeks to decentralize and distribute access to data.

Value-Based Care/Engagement

No longer are patients considered passive recipients of care. Healthcare is now obliged to engage patients in their health, healthcare decision-making, care, and treatment. Engagement can be maintained through digital means. Note that patient engagement is not to be confused with patient experience, which is the pathway (journey) a patient may take.

Financial drivers have already seen healthcare practices becoming increasingly incentivized to best engage with each patient, to ensure services received are satisfactory and of quality.

One of the key drivers of data-driven solutions is the demand for patient engagement and the transition toward value-based care. Better patient engagement enhances trust between patients, treatment providers, and bill payers. Moreover, it leads to better health outcomes and cost savings (or some other benefit) to the provider.

Pioneering health insurance initiatives seek to engage patients with the goal of better health outcomes through entwining premiums to good health. Health innovators Diabetes Digital Media in Warwick, England,

are working with global insurance providers and health organizations to provide scalable digital health technologies for users to optimize their health and wellbeing.[27]

Blue Shield of California is improving patient outcomes by developing an integrated system that connects doctors, hospitals, and health coverage to the patient's broader health data to deliver evidence-based, personalized care.[28] The aim is to help improve performance in disease prevention and care coordination.

Healthcare IoT—Real-Time Notifications, Alerts, Automation

Millions of people use devices that datafy their lives toward the quantified self. Devices connected to the Internet currently include weighing scales; activity monitors (such as Fitbit, Apple Watch, Microsoft Band) that measure heart rate, movement, and sleep; and blood glucose meters: all of which send metrics in real time and track user behavior in up-to-the-second fashion. Early 2018 even saw the first fetal heartbeat wearable.[29] The data recorded could be used to detect the risk of disease, alert doctors, or request emergency services depending on the biometrics received. Many integrated devices are going beyond pulse and movement to measure sweat, oxidation levels, blood glucose, nicotine consumption, and more.

With sophisticated devices come sophisticated solutions to novel problems. For instance, now that heart rate monitoring is cheaper and more pervasive, conditions such as AFIB can be detected far easier and far earlier than ever before. Fluttering of one's heart rate above 300 bpm (rather than 60–80 bpm) could be a symptom of AFIB. Patients diagnosed

with the condition are 33% more likely to develop dementia; and more than 70% of patients will die from a stroke.[30] The treatment can often be simple: anticoagulants and blood thinners, which are up to 80% effective. Healthcare providers are moving slowly toward sophisticated toolkits to utilize the massive data stream created by patients and to react every time the results appear disturbing. This adoption is being driven by established digital organizations and startups in conjunction with health bill payers.

In another use case, an innovative program from the University of California, Irvine, gave patients with heart disease the opportunity to return home with a wireless weighing scale and weigh themselves at regular intervals. Predictive analytics algorithms determined unsafe weight gain thresholds and alerted physicians to see the patient proactively before an emergency readmittance was necessary.[31]

What is interesting to recognize is that these connected health devices do not necessarily motivate the audiences or populations most at risk of adverse health events. According to several randomized trials, Fitbit wearers do exercise more, but not enough to guarantee weight loss and improved fitness. In fact, some studies have determined they can be demotivating.[32] Then there is the question of accuracy. A Cleveland Clinic study in 2016 found that heart rate monitors from four brands on the market were reporting inaccurate readings 10 to 20% of the time, which demonstrates that there is some way to go in the precision of the technology.[33]

While these devices show potential, there is still the problem of an average abandonment rate of more than 30% after a period of engaged usage. The software application is now the channel for engagement, used as an intelligent layer on top of devices and IoT, where they can also introduce behavior change psychology for sustainable behavior change and usage.

Offering users of devices such as these tangible incentives like discounts on health or life insurance could become more mainstream,

and drive prevention of many chronic lifestyle-related diseases similar to incentivized car insurance with embedded black-box sensors.

Real-time alerting can also be used to notify patients of adverse effects of medications they are prescribed. Currently, patients wouldn't receive any notification unless they were registered with the treatment provider. Healthcare providers can use public feeds to notify their patients of potential adverse effects. E-mails or text messages would suffice as a form of engagement and could provide as much instruction as required, reducing time with the clinician.

EHRs can also trigger warnings, alerts, and reminders for when a patient should get a new lab test; or track prescriptions to see if a patient has been following treatment instructions.

Movement Toward Evidence-Based Medicine

Evidence-based medicine is a term denoting treatment given based on proven scientific methods in pursuit of the best possible outcomes. Clinical trials work on a small scale, testing new treatments in small groups with internal validity (i.e., no other conditions or concerns other than those specified), and looking at how well treatments work and establishing if there are any side effects. With growing datafication, there is also increasing “real-world evidence” or data, which can be analyzed at an individual level to create a patient data model and aggregated across populations to derive larger insights around disease prevalence, treatment, engagement, and outcomes. This approach improves quality of care, transparency, outcomes, value, and at its core, democratizes healthcare delivery.

Using data for similar groups of patients, we can understand the treatment plans patients with similar profiles have taken, allowing the best treatment plan to be recommended based on how others in the population responded to any given treatment. The recommended treatment pathway would be best for the patient and could be explained to both the patient

and healthcare professional as to why it was the best pathway to follow. This has roots in predictive analytics and goes beyond identification of populations and treatments.

By mining real-world patient data, including clinical patient records and real-time, personal, patient health data—and with demographic and health data on disease prevalence, treatment pathways, and outcomes—we can learn about optimal treatment plans for individuals, facilitating precision medicine, which is the pinnacle of evidence-based medicine.

By connecting real-world patient and clinical data to the genome, it can be personalized down to the genetic makeup of patients and populations for the following:

- Prescribing of medications
- Adverse effects and reactions
- Prevention strategies
- Likelihood risk of future diseases

Public Health

Analysis of disease patterns and outbreaks allows public health to be substantially improved through an analytics-driven approach. Big data can assist in determining needs, required services, or treatments; and it can predict and prevent future crises to benefit the population. By mapping patient location, it would be possible to predict outbreaks, such as influenza, that could spread within an area, making it easier to formulate plans for dealing with patients, vaccinations, and care delivery.

In West Africa, mobile phone location data proved invaluable in tracking the spread of the population—and as a result, helped to predict the Ebola virus's expanse.[34]

After the Haiti earthquake in 2010, a team from Karolinska Institute in Sweden and Columbia University in the United States analyzed calling

data from two million mobile phones on the Digicel Haiti network.[35] Phone records were used to understand population movements and for the United Nations to allocate resources more efficiently. The data was also used to identify areas at risk of the subsequent cholera outbreak.

Evolution of Data and Its Analytics

Proximity is driving the future. As people, patients or agents in a digital world, we expect a relationship regarding nearness in time between data that is given and feedback that is received. As the scale of data increases, so too does the demand that the time between data creation to insight and action be reduced. There is substantial economic value in reducing the time from detected event to automatic response. This could help spot fraudulent transactions before a transaction completes, or even serve information matched to an interest or need.

Just as data has evolved, so too has the analytics of data (Figure 2-5). As a concept, the analysis of data, or analytics, has gone through several iterations, driven by the demand for access to data to drive decision-making.

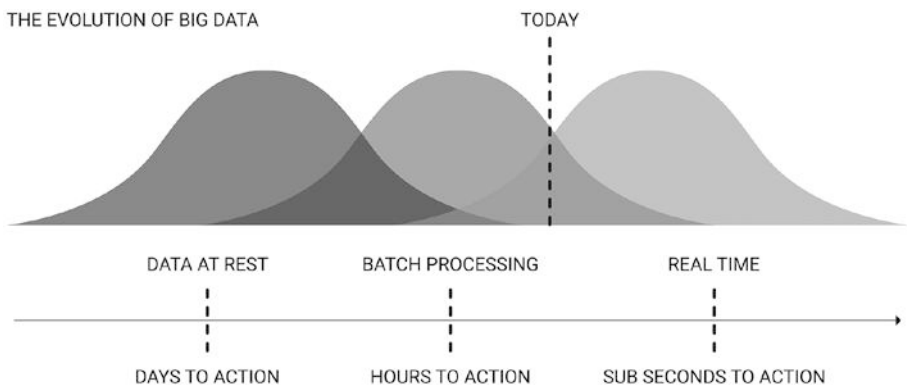


Figure 2-5. *The evolution of big data and its analytics*

Analytics 1.0 took the form of traditional analytics on small, slower velocity datasets, often in the form of spreadsheets and static documents. This is suitable for little data. Typical analytics were descriptive and report-based. Data would usually be structured in nature. This drove business intelligence in the 1990s with predefined queries and detailed or historic views—typically driven by web dashboards leveraging structured data like customer data, behavioral data, sales data, patient records, and so forth. Data about production processes, sales, interactions, and transactions were collected, aggregated, and analyzed within traditional relational databases. In the Analytics 1.0 era, data scientists spent much of their time preparing data for analysis rather than performing analytics themselves.

Analytics 2.0 marked the integration of big data with traditional analytics, with interfaces to support the querying of big datasets in real time. Big data analytics began in the 2000s, which facilitated competitive insight through complex queries combined with predictive views leveraging both structured and unstructured data such as social media, behavioral data, and the growing variety of user data. Big data that couldn't fit or be analyzed fast enough on a centralized platform would be processed with frameworks such as Hadoop, an open source software framework for fast batch data processing across parallel servers both in the cloud or on-premise.

Unstructured data would be supported through more than just SQL (Structured Query Language) databases, supporting key and value, document, graph, columnar, and geospatial data. Other big data technologies introduced during this period include “in memory” analytics for fast analysis whereby data is managed and processed in memory rather than on disk.

Today, vast amounts of data are being created at the edge of the network, and the traditional ways of doing analytics are no longer viable. Organizations that engage with consumers—whether they make things, move things, consume things, or work with customers—have increasing

amounts of data on those devices and activities. Each device, shipment, and person leave a trail referred to as a data exhaust.

Analytics 3.0 marked the stage of maturity in which organizations realized measurable business impact from the combination of traditional analytics and big data. The approach supports agile insight discovery and real-time analysis at the point of decision. Analytics 3.0 is mostly a combination of conventional business intelligence, big data, and the IoT distributed throughout the network. Organizations can analyze those datasets for the benefit of customers and monetize them. They also can embed advanced analytics and optimization in near real time into every product or service decision made on the front lines of their operations.

Turning Data into Information: Using Big Data

Data's value is demonstrated in the ability to convert it into information that can drive actionable insight to drive behaviors and workflows. Data analysis/analytics is a subset of data science that seeks to draw insights from sources of raw data. Today's advances in analyzing big data allow researchers to decode human DNA in minutes; predict where terrorists plan to attack; determine which gene is most likely to be responsible for specific diseases; and, of course, which ads you are most likely to respond to on Facebook.

Analysis of big data can help organizations understand the information contained in their data as well as identifying data most important to business objectives, outcomes, and decisions. Stakeholders typically want the knowledge that comes from analyzing data, which is determined based on the type of results produced and the reasoning approach.

The majority of big data is unstructured. IBM estimates that 80% of the world's big data is unstructured, which includes photos, audio, video, documents, and interactions.[36]

Data needs to be processed to be of value. For this to take place, resources are required—in the form of time, talent, and money. Herein is one of the inertias toward adopting big data approaches: is processing the data into something with usable insight worth the capital and risk of doing so?

There are four streams within the data science analytics taxonomy (Figure 2-6). The classifications are based on the types of results that are produced.

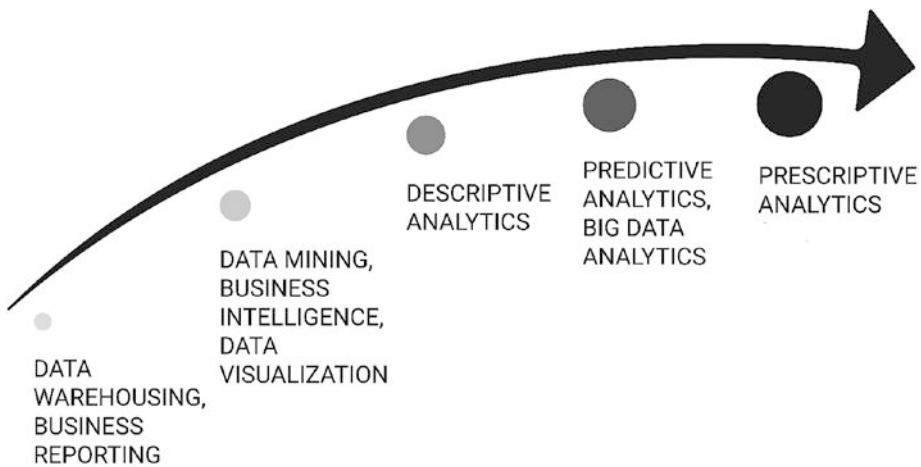


Figure 2-6. *Streams of analytics*

Descriptive Analytics

Descriptive analytics brings insight to the past, focussing on the question of what happened, through analyzing data from history. Descriptive analytics uses techniques such as data aggregation and data mining to provide historic understanding. There is much to be learned from descriptive analytics. Common examples of descriptive analytics are reports that provide the answers to questions such as how many patients were admitted to a hospital last July, how many were readmitted within 30

days, and how many caught an infection or suffered from a mistake in care. Descriptive analytics serves as the first form of big data analysis.

This basic level of reporting is still lacking for many organizations. Proprietary data standards maintain unhelpful pits of information; a lack of available human expertise or proper organizational buy-in can leave data collecting and unused.

Data can become more valuable through dataset linking. For instance, linking physician appointments to hospital admissions, to education attendance, medication adherence, and access to services can provide substantial insight and understanding for health delivery optimization and cost savings.

Limitations of descriptive analytics are that it gives limited ability to guide decisions because it is based on a snapshot of the past. Although this is useful, it is not always indicative of the future.

Diagnostic Analytics

Diagnostic analytics is a form of analytics that examines data to answer the question of why something happened. Diagnostic analytics comprises of techniques such as decision trees, data discovery, data mining, and correlations.

The next two types of analytics, predictive and prescriptive analytics, require the use of learning algorithms and are covered in more depth in Chapter 3.

Predictive Analytics

Predictive analytics allows us to understand the future and predict the likelihood of a future outcome. Predictive analytics uses the data that you have at your disposal and attempts to fill in missing data with best-guess estimates.

Predictive analytics is characterized by techniques such as regression analysis, multivariate statistics, data mining, pattern matching, predictive modeling, and machine learning. Predictive analytics uses historical and current data to forecast the likelihood of future events or actionable outcomes.

Predictive analytics skills are in demand as healthcare providers seek evidence-based ways to reduce costs, take advantage of value-based reimbursements, and avoid the penalties associated with the failure to control chronic diseases and adverse events that are within their power to prevent. Predictive analytics is in the midst of disruption. The last 5 years of innovation has made significant progress with measurable impacts to the lives of patients. Wearable technology and mobile apps now make it possible to spot conditions such as asthma, atrial fibrillation and COPD.[37]

Predictive analytics remains elusive, as it requires access to real-time data that allows near real-time clinical decision-making. To support this, medical sensors and connected devices must be fully integrated to provide up-to-the-moment information on patient health. Besides, clinicians need to be experienced in using such data. As well as using the individual patient's data, more accurate diagnoses and treatments must draw from as much patient data as is available, including population-level data.

As data sources and technologies develop, advanced decision support from cognitive computing engines, natural language processing, and predictive analytics can help providers pinpoint diagnoses that may otherwise elude them. Population health management tools can highlight those most at risk of hospital readmission, at risk of developing costly chronic diseases, or responding adversely to medication.

Use Case: Realizing Personalized Care

Metabolic health has been demonstrated to influence the risk of a variety of health conditions, including type 2 diabetes, hypertension, some dementias, and some cancers. While acting as a digital intervention for type 2 diabetes, the Low Carb Program app used a variety of features in the form of health biomarkers and demographics including blood glucose, weight, gender, and ethnicity to score a patient's metabolic health. As the app expanded to integrate with wearable health devices, increased weight and blood glucose data enabled the algorithm to extend to determine the likely risk of pancreatic cancer. This is communicated to patients with their healthcare team where relevant and is referenced through cohort data comparison.[38]

Users receive supported notifications that their data does not seem to fit within the demographic that is expected of them and prompts the user to speak to their healthcare professional. As demonstrated in this example, as further biosensors and algorithms are developed, the ethical implications of predictive analytics is significant.

Use Case: Patient Monitoring in Real Time

A typical hospital ward would see ward nurses manually visit patients to monitor and confirm vital health signs. However, there is no guarantee that the patient's health status won't deviate, for better or worse, between the time of scheduled visits. As a result, caregivers often react to problems as the result of adverse events, whereas arriving earlier in the process can make a significant difference to patient well-being. Wireless sensors can capture and transmit patient vitals far more frequently than human beings can visit the bedside. Alan Gates' Hortonworks used this data to provide caregivers data in real time to enable them to respond more promptly to unexpected changes.

Over time, the same data could be applied to predictive analytic techniques to determine the likelihood of an emergency before it can be detected with a bedside visit.

Prescriptive Analytics

Prescriptive analytics strives to make decisions for optimal outcomes: that is, to use all available data and analytics to inform and evolve a decision of what action to take—that is, smarter decisions. Prescriptive analytics attempt to quantify the effect of future decisions to advise on possible outcomes before decisions are made. At its best, prescriptive analytics predicts not only what will happen but also why it will happen to provide recommendations regarding actions that will take advantage of the predictions.

Prescriptive analytics use a combination of AI techniques and tools such as machine learning, data mining, and computational modeling procedures. These techniques are applied against input from many different datasets including historical and transactional data, real-time data feeds, and big datasets.

Prescriptive analytics performs an in-context prediction, taking into account the available evidence. For example, typical hospital readmission analytics display a forecast of patients likely to return in the next 14 to 30 days. A more useful predictor may integrate the same dashboard with projected costs, real-time bed counts, available education material, and follow-up care. With additional information, clinicians can determine patients with a high risk of readmission and take actions to minimize this risk, resulting in improved patient outcomes and reduced consumption of hospital resources. Population management approaches could consider patients who are obese, adding filters for factors such as type 2 diabetes risk and metabolic syndrome risk to determine where to focus treatment and which treatments to target. Pharmaceutical companies can use prescriptive analytics to expedite drug development through identifying

audiences that are most suitable for clinical trials. This includes patients presumed to be compliant and those not expected to have complications.

Prescriptive analytics enables the optimization of population health management through identification of appropriate intervention models for individuals from risk stratified populations—combining clinical, patient, and wider available health data.

Analytics is undergoing a further transformation, from prescriptive analytics through to contextual analytics. Contextual analytics considers wider data—such as environmental, location, and situational data—when prescribing the best course of action.

Use Case: From Digital To Pharmacology

Different treatment pathways work differently for different people. Collected on a large scale, physicians can prescribe therapies based on the latest evidence-based and predictive and prescriptive analyses of patient populations. Incorporating digital education as well as pharmacological therapies, delivery of treatments through such facilities democratizes healthcare access by ensuring the most precise treatment is always prescribed, alongside ensuring resources are focused where required.

Reasoning

A system can determine conclusions with information available in a knowledge base using three main methods: deduction, induction, and abduction. Although a background in logic is not required, it is useful to understand the differences between the different reasoning approaches. Any data scientist should be familiar with these reasoning techniques.

Deduction

Deductive reasoning allows you to make statements that are necessitated by facts that you know. For example, you are given two facts:

1. It rains every Saturday.
2. Today is Saturday.

Deductive reasoning allows you to determine the true statement that it is going to rain today if today is Saturday.

In deductive reasoning, one infers a proposition q , which is logically sensical from a premise p . Most reporting systems and business intelligence software is deductive.

Induction

Inductive reasoning enables you to make statements based on the evidence you have accumulated until now. However, the key point is that evidence is not the same as fact. Substantial evidence for something only suggests, however strongly, that some “thing” is a fact. Consequently, statements that are determined are only very likely to be true in an inductive approach, rather than absolute.

With inductive reasoning, one attempts to infer a proposition q , which follows from p , but which is not necessarily a logical consequence of p .

For example, if it has rained in Coventry, England, during December for a recorded history of over 50 years, you would have strong evidence (data) for the inductive statement that it will rain in the coming December in Coventry. However, this does not mean it is going to happen; it is not a fact.

Statistical learning is about inductive reasoning: looking at some data, scientifically guessing at a general hypothesis, and making statements or predictions on test data based on this premise.

Abduction

Abduction is an adaptation of inductive reasoning. Abductive reasoning attempts to use a hypothesis p to explain a proposition q . In abduction, reasoning flows in the opposite direction to deductive reasoning. The best hypothesis, that is, the one that most effectively explains the data, is inferred to be the most probably correct one.

A classic example is the following:

- (q) Observation: When I woke up, the grass outside my window was wet.
- (p) Knowledge base includes information: Rain can make the grass wet.
- Abductive inference: It probably rained at night.

There are often many up to infinite possible outcomes, so abductive reasoning is useful in determining how to prioritize explanations to investigate. Big datasets provide opportunities for induction and inference through machine learning.

How Much Data Do I Need for My Project?

Although big data is defined by the fact it is stored in distributed systems, when it comes to creating machine learning algorithms, it's about quality and not quantity. For instance, if you have over 1,000 data points for a patient, how are you to understand the most important attributes from what is noise? Further still, is it likely that you would get better outcomes from 5,000, 10,000, or 100,000 data points? This problem is best addressed through data science and is entirely dependent on the domain and the quality of your data samples.

The future of evidence-based medicine is the realization of data-driven, personalized healthcare. Data-driven healthcare occurs when

there is a synergetic partnership between healthcare providers, patients, and data obtained from all clinical documentation. This results in analytics that drives actionable insights in real time, allowing for an enhanced patient experience and streamlined clinical and administrative workflows.

Challenges of Big Data

There are several challenges in the development of big data initiatives.

Data Growth

By the very nature of big data, storage is a challenge. There is broad agreement that the size of the digital universe will double every 2 years, which equates to a fiftyfold growth from 2010 to 2020. Most big data is unstructured, meaning that it doesn't reside in a traditional database schema. Documents, photos, audio, videos, and other unstructured data can be difficult to search, analyze, and retrieve. Hence, management of unstructured data is a growing challenge. Big data projects can evolve as quickly as the data used within them.

Infrastructure

Big data consumes technical resources in the form of infrastructure, storage, bandwidth, databases, and so forth. The challenge is not technical so much as one of finding reliable vendors of services and support, and of getting the correct economic model of remuneration. A solution delivered from the cloud will have greater scalability, cost-effectiveness, and efficiency compared to an on-premises solution.

Expertise

Finding good quality expertise (or a deep stack expertise) in the disciplines of data analysis and data science is a challenge faced by many organizations. Compared to the amount of data generated, there are very few data scientists in proportion.

Data Sources

A significant challenge with big data is the volume and velocity at which data is generated and delivered. Managing the plethora of incoming data sources is a task in itself.

Quality of Data

Data quality is certainly not a new concern, but the ability to store each piece of data generated in real time compounds the problem. Common causes of dirty data must be addressed and include user input errors, duplicate data, and incorrect data linking. As well as maintaining and cleaning data, big data algorithms can also be used to help clean data.

Security

There are significant risks in the security and privacy of patient data. With sensitive health data, there is rightly an increased sensitivity to security and privacy concerns. The tools used for analysis, storage, management, and utilization of data are from a variety of sources, which exposes the data to risk. Cybersecurity is a concern. Significant data breaches have occurred in the healthcare industry. In the United Kingdom, the Information Commissioner (ICO) fined the Brighton and Sussex University Hospitals NHS Trust[39] after it was found that the sensitive data of several thousand people were resident on hard drives placed up for auction on

eBay. To date, the most substantial healthcare breach recorded has been of American health insurance company Anthem.[40] The breach exposed the personal records, including home addresses and personal information, of over 70 million current and former members. Data loss, including data theft, is a genuine problem that is continuously tested to its limits.

Security challenges include authentication for every team and team member accessing datasets, restricting access based on user needs, recording data access histories, and meeting compliance regulations and ensuring data is encrypted to protect from snoopers and any malintent.

Security of medical devices poses a unique threat because of their technological diversity. Medical devices, from health applications on a smartphone to insulin pumps, are becoming increasingly networked, leaving unique openings for hackers.[41] This kind of activity is regarded as dangerous, although for those more technical, it is understood that this can be utilized for positive benefit. Insulin pumps and continuous glucose monitors have been “hacked” to function as an artificial pancreas by enthusiasts way before pharmaceutical company or digital companies had. However, the key concern is that, if exploited, similar device vulnerabilities could lead not only to data breaches but fatalities in people relying on medical devices. Today, the balance between disruption and progress is one of real-world experience.

Then there is the topic of third-party services and vendors. Most health devices and health apps offer API access. It’s useful to remember that external vendor services and APIs are only as good as the developers that have developed them.

Resistance

Lack of understanding of value and outcomes from machine learning projects can leave internal stakeholder resistance. You can only manage a project efficiently if you can ascribe clear value to results and prioritize

clear explanations of outputs to stakeholders. Over and above inertia, once key patterns have been identified, organizations must be prepared to make necessary actions and implement changes to derive value from them. This is often a primary inertia, so precise and jargon-free documentation and governance is useful for understanding the project and governance purposes.

Policies and Governance

The vast amount of big data captured through wearable devices and sensors raises questions as to how data is collected, how data is treated, how data is analyzed, and how data should affect the policy-making process. Besides, data privacy and security are legal areas in a state of constant development and evolution.

Fragmentation

Most organizational data is highly fragmented. Hospital teams have their patient data, as do physicians and tertiary care teams. This creates challenges at several levels: syntactic (i.e., defining common formats across teams, sites, and organizations), semantic (agreeing on common definitions), and political (determining and confirming ownership and responsibilities). This is before the patient's own data, collected through phones, health apps, nutrition apps, and so forth.

Lack of Data Strategy

To benefit from data science investments, organizations must have a coherent data strategy. It's key to establish the need for the data and identify the right type and sources of data for use. What is the desired goal with the data and what are the objectives? It is important to establish the data needs, to save wasting resources and efforts collecting data that are of no use.

Visualization

To truly benefit from data science investments, stakeholders must understand the outputs of data analysis and exploration. Poor visualization of data is connected to lack of data strategy: if specific use cases and objectives are not determined, visualization of results from data can prove challenging due to a lack of clear direction.

Timeliness of Analysis

Linked to proximity, the value of data may decrease over time. Applications such as fraud detection in banking require data to be delivered as close to real-time as possible to be effective.

Ethics

Data ethics is emerging as a topic as datasets become big enough to raise practical rather than theoretical concerns about ethics. Patients are leaving a constant data trail through sensors, wearables, transactions, social media, and transportation. For instance, the genetic testing company 23andMe can sell to individuals reports on their genetic risk for ten diseases, including Alzheimer's disease and Parkinson's disease. This leads to challenging conversations about what is and what is not appropriate, discussed more in Chapter 8.

Data and Information Governance

Data is a precious asset to any business or organization; but in healthcare, data governance is paramount to ensure consistency, safety, and progression of any organization toward becoming data-driven and analytically driven. The need for data governance is driven primarily by

the requirement for accountability in a risk-averse industry. As the value of data is realized, data governance too is evolving to define the approach, management, and utilization of data in an organization. Processes and standards such as ISO (International Organization for Standardization) 10027 and GDPR (General Data Protection Regulation) have been used to bring the entire data industry in line with best practice.

Data governance and information governance are often used interchangeably but do differ.

Governance of data refers to the practice of managing data assets throughout their life cycle to ensure that they meet standards of quality and integrity. The goal of data governance is to provide user trust in the data through optimal data validity and accountability. It sets the framework for creating high-quality data and using data in a secure, ethical, and consensual manner. Data governance in healthcare seeks to improve efficiency, create accountability, and establish a pool of data that providers and patients can use to make proper health decisions.

Often the best approach to data governance is to govern to the least extent necessary to achieve the most common good. In several cases, governance is being set for circumstances that don't require governing yet. Organizations can waste tremendous amounts of time through management and decision-making, including constraints on data, which can hinder project development and direction. Healthcare organizations should be motivated by the acquisition of more data to learn—to gather information to manage risks better and improve outcomes.

Data governance is required to uphold a variety of criteria to return quality information.

Data Stewardship

The steward is accountable for the data as defined and its appropriate use. The objective of this aspect of data governance is to assure quality regarding data accuracy; data accessibility; and data consistency,

completeness, and updating. This is typically the organization running the project. Further still, most organizations would have an appointed officer(s), responsible for accessing the data, with elevated privileges.

Data Quality

Ensuring data quality is arguably the most important function of data governance besides security. Poor data quality has a detrimental impact on accuracy, particularly for learning projects. Governance should seek to enforce quality, accuracy, and timeliness of data.

Data Security

Security of data is essential. First and foremost, data governance sets out how the data is protected, typically how it is encrypted, who has access to it, processes for handling data—including how the data may be used—and what is done in case of a breach. This is motivated by several drivers: regulatory requirements that demand action (such as GDPR taking place in Europe), the growing risk of cybersecurity, and the fact that perceptions of digital weakness can harm reputation and trust. This is demonstrated in how the number of web sites that operate with a secure <https://> protocol has risen by 4,000% since 2012.[42] Patients expect their data to be secure.

Data Availability

Fragmentation of data is one of the biggest problems healthcare organizations face. Ensuring stakeholders' providers' and patients' data access is critical. Patients "own" their data; and as a result, access needs to be transparent, quick, and efficient for all. Data access governance may set permissions and authentication levels for user roles to systems.

Data Content

The data and information governance for a project may state the types of data collected, including health data; metadata; and location, profile, and behavior data. Data governance may also detail what the data is used for.

Master Data Management (MDM)

As data transfer becomes a priority, the MDM acts as a master reference to ensure consistent use of data across organizations. The goal of MDM is to create a common point of reference that can be shared within an organization.

Use Cases

In almost all cases, data governance policies will include use cases for data access and “what would happen if” scenarios. It is useful to have all team members required in action to be involved in dummy runs to ensure that data access or access requests are not met with periods of silence.

Use cases would include the set of processes that would take place in light of a request for data from a patient, request for deletion, ensuring appropriate roles were able to access the data, ensuring appropriate transactional logs were kept, and so forth.

Data governance provides several benefits over and above those that are internal, including the following:

- Protecting the interests of data stakeholders—particularly the data “giver”
- Standardizing procedures and processes for streamlined repetition and minimization of error

- Reducing costs and improving effectiveness
- Greater transparency and accountability between data transfer parties

Information governance has a slightly different purpose. Information governance is the management and control of information. This information is formed through the use of data assets (Figure 2-7). A single data point, such as a blood glucose reading, without any additional context or metadata offers very little. However, a series of blood glucose readings over the past 3 months can show whether the patient has good or bad blood glucose control, for example, which can be a risk factor for comorbidities. It pays attention to personal and sensitive information linked to patients and employees.



Figure 2-7. *Data vs. information—what's the difference?*

In most organizations, data and information governance is consolidated; but it is useful to understand the distinction.

A resolute data governance program would involve a governing body or council, a defined set of procedures, and the way in which said procedures are upheld and verified.

As learning systems develop, information governance is becoming more critical. Digital health apps and integrated bio-human health devices will soon be able to predict and determine risk and time to disease as datasets grow infinitely large. Regardless of the question of accuracy, legislation on what can and can't happen with this information is required.

Deploying a Big Data Project

Getting started with any big data project involves three steps.

1. Understand how the project will impact your organization

This involves identifying the business case; objectives (what is it I want to achieve?); current and required infrastructures; data sources; and quality, tools, and KPIs regarding measuring success. Identifying quantitative objectives is key to measure success and performance. Determining a measurable ROI to determine the success of the project is useful to demonstrate value to stakeholders. Identifying goals with members from technical, medical, and operational teams help to maintain clarity of focus and direction. Identify potential use cases.

2. Find the skills and technology you need

Identification of the use cases for the project can help define the technology, infrastructure, and capabilities required for the project. Skills in Python and R (open source programming languages) are most likely going to be required, with analytics

skills also desired for data science. This can be outsourced, contracted in, or developed into full-time roles.

3. Implement the project

Identify the data you wish to include and make a job of identifying data that will also be excluded from use cases. Identify the types of analyses required on the data to gather the desired output from the project. Explore the requirements necessary for collecting and preparing data into usable formats for algorithms; the governance required for holding and using the data; and how it will be presented to stakeholders and, most importantly, the end users. Architectural decisions such as database structures, platform providers, and type of models may be made at this point. It's at this point you can consider the types of analytics applied to the data.

You may wish to consider phasing of the project and the gaps between current and future capabilities that require addressing.

Develop the project in a test environment and present results to users in a meaningful way.

Big Data Tools

There are various open source and paid-for tools that can be used for big data collection and machine learning projects. The fundamentals begin with Hadoop and NoSQL (non-relational) databases.

Hadoop is synonymous with big data, and is a Java-based, open source software framework, or software ecosystem, for distributed storage of large datasets on computer clusters. Hadoop can process big

and complex datasets for insights and answers. Hadoop Distributed File System (HDFS) is the storage system used by Hadoop that separates big data and distributes it across nodes within a cluster. This can also replicate data in a cluster, providing high availability. For those not wanting to delve into Java themselves, Microsoft Azure, Cloudera, Google Cloud Platform, and Amazon Web Services all provide big data hosting services powered by Hadoop.

MongoDB is a modern approach to databases that is useful for managing data that changes frequently, or data that is unstructured or semi-structured in nature. The architecture has better write performance, takes less storage space with compression, and reduces operational overhead. MongoDB is part of a growing movement known as NoSQL databases, which are used to store unstructured data with no particular schema. Rows can all have their own set of unique column values. NoSQL has been driven by the requirements to provide better performance in storing big data.

R is an open source software and programming language used for data mining and statistical analysis and graphical techniques.

Python is another popular open source programming language used in machine learning and data mining for data preprocessing, computing, and modeling.

Conclusion

Big data and healthcare mean significant progress toward realizing a plethora of opportunities. This includes remote patient monitoring, precision medicine, preventing readmissions, and advancing disease understanding. From mundane tasks of data entry and record keeping to advanced and functional applications like observing and understanding people's blood glucose patterns—data is becoming part of our day-to-day fabric.

The potential for positive impact that big data-driven solutions can bring seemingly outweighs any negative sentiment. The growing trend toward centralization of medical data, toward the custodianship of the patient, is causing concern; yet as long as privacy and security are maintained, it is likely to facilitate the development of new treatments and contribute toward the evolution of medical understanding. The healthcare industry remains well within its infancy of leveraging big data for clinical and business use. The explosion of wearables has allowed the measurement of heart rate or steps to be taken during a day, which may seem like a relatively ordinary dataset; but the potential impact on preventative medicine and the general health of a population is immense.

Big data provides the opportunity for precision medicine through the delivery of data-driven, evidence-based medicine. The benefits are plentiful: first and foremost through the delivery of more precise treatments for patients. With this comes cost savings through fewer mistakes, fewer hospital admissions, decreased mortality, and improved patient satisfaction. This also saves money by optimizing the use of the physician and wider healthcare team's time, as more time is spent with patients in need.

CHAPTER 3

What Is Machine Learning?

The first Checkers program from machine learning pioneer Arthur Samuel debuted in 1956, demonstrating artificial “intelligence” capabilities.[43] Since then, not only has the application of AI grown, there is now a velocity, volume, and variety of data that has never been seen before. Samuel’s software ran on the IBM 701, a computer the size of a double bed. Data was typically discrete. Almost 70 years later, data is ubiquitous, and computers now more powerful such that 100 IBM 701s can fit into the palm of our hand. This has facilitated the subsets of machine learning and deep learning within AI (see Figure 3-1).

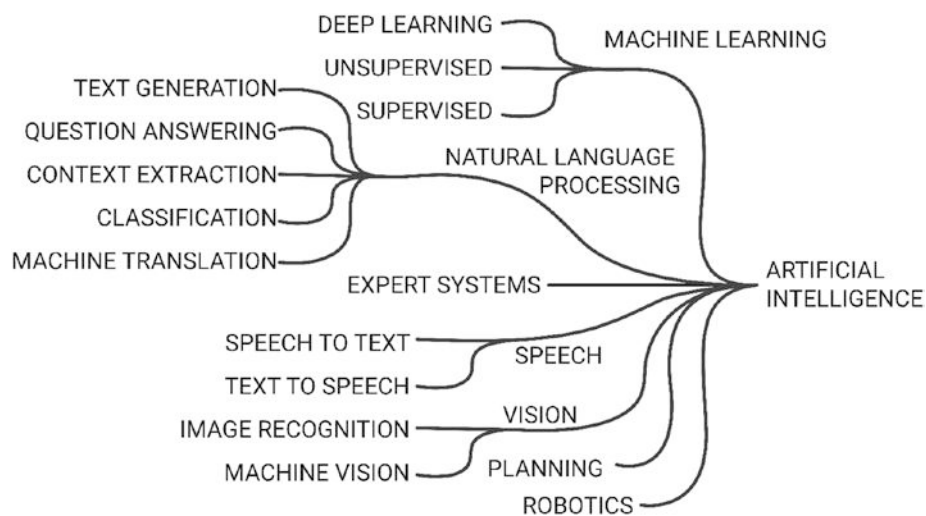


Figure 3-1. *Components of AI*

We see and engage with “artificial” intelligence, whether developed by machine learning or otherwise, in our everyday lives: in ordering an Uber, ordering food, looking for an Airbnb, searching on eBay, asking Alexa a question, and in digital education programs. Pondering on the AI humans are able to develop is limited only by the human imagination. New and innovative devices are increasingly connected to the cloud and increasing the wealth of opportunities in healthcare. With a wealth of data at its helm as an industry, AI is shifting its focus onto the power of machine learning and deep learning, where machines can learn for themselves—and this is accelerating current-generation breakthroughs in the AI space. If Samuel’s Checkers program demonstrated simple AI, machine learning is best demonstrated through AlphaGo Zero, which learned to play checkers simply by playing games against itself, surpassed human level of play, and defeated the previous version of AlphaGo by 100 games to 0.

Machine learning systems are not uncommon in healthcare. Researchers at Houston Methodist Research Institute in Texas have developed an agent that is enabling review and translation of mammograms 30 times faster than a human doctor and with 99% accuracy.[44] According to the American Cancer Society, 50% of the 12.1 million mammograms performed annually yield false results, meaning that healthy women get told they have cancer.[45] The software, which can be understood as an autonomous entity, learned from millions of mammogram records, reducing the risk of false positives.

There are many books and resources on the field of AI and machine learning. Although a comprehensive and mathematically grounded discipline, this book aims to cover the topic for the discerned reader, without delving too deep into the semantics of linear algebra, probability, and statistics.

This chapter will cover many of the topics within machine learning, beginning with an introduction to the basics of AI and terminologies and then moving on to artificial neural networks, perceptrons, natural language processing, and expert systems. Algorithms are covered in Chapter 4.

Basics

Before delving into the depths of AI and machine learning, it is useful to understand several key definitions. A brief guide to useful terminology is included in the following, with a more in-depth glossary at the end of this book.

Agent

An agent is anything that can be seen as perceiving its environment through sensors and acting on them through effectors. This takes Russell and Norvig's approach to agents. As a result, this includes robots, humans, and software programs. There are many different types of agents, which

can be classified into categories based on their characteristics. *Artificial Intelligence: A Modern Approach* (AIMA) is a comprehensive textbook on artificial intelligence written by AI experts Stuart J. Russell and Peter Norvig.[46]

How should the agent behave? Rational agents are agents that do the right thing. This can be represented by a mathematical measure, or evaluation of the agent's performance, which the agent can seek to maximize.

An agent can be understood as a mapping between percept sequences and actions. For example, a medical diagnosis system may treat a patient in a hospital as its environment: perceiving symptoms, findings, and patient's answers. It may then act on these through questions, tests, and treatments with the goal of having a healthy patient. The definition of an agent applies to a software program or algorithm.

Autonomy

In AI, autonomy refers to the extent to which the agent's behaviors are determined by its own experience. If there is no autonomy, all decisions are made from a static dataset or knowledge base. In a completely autonomous situation, an agent would make decisions at random. Agents have the explicit ability to become more autonomous over time. Software, just like humans, can learn from experiences over time. Systems can learn from experiences in their environment (at its most basal, "learning" from some new data) to become more autonomous in it.

Interface

An interface agent is one that provides an interface to a complex system.

Performance

Performance is the measure used to evaluate the behavior of the agent in the environment. It answers the question, Does the agent do what it's supposed to do in the environment?

Goals

Goals refer to what the agent is trying to achieve. In the field of machine learning, the goal is not to be able to get the agent to “learn.” Instead, this refers to the overall goal or objective.

Goals can be broken down into smaller tasks. For example, to accomplish the goal of a Caesarean section, the tasks could be considered cutting open the abdomen, bringing the baby out, checking the baby is healthy, and stitching the abdomen back up again.

Utility

Utility refers to the agent's own, internal performance assessment—that is, the agent's own measure of performance at any given state. This may differ from the performance of the agent. It is noteworthy that the utility function is not necessarily the same as the performance measure, although it can be. An agent may have no explicit utility function, whereas there is always some performance measure. A utility function is used to map a state to a real number, which describes the associated degree of happiness or progress toward a goal or objective. This allows rational decisions in cases where there are several paths to the same goal, allowing an agent to distinguish the paths better than others.

Knowledge

Knowledge is acquired by an agent through its sensors or knowledge about the environment. Knowledge can be used to decide how to act; if it is stored, it can be used to store previous knowledge states (or in other words, a history); and it can be used to determine how actions may affect the environment.

Machine learning is data-driven. Thus, the dawn of ubiquitous data and pervasive computing is fantastic for developing systems with real-world application.

Environment

The environment is the state of the world for an agent. An environment has several characteristics, following Russell and Norvig's lead.

- Accessibility

This refers to how much of the environment the agent has access to. In cases of missing information, agents may need to make informed guesses to act rationally.

- Determinism

Environments can be deterministic or nondeterministic. Deterministic environments are environments where the exact state of the world can be understood. In this instance, utility functions can understand their performance. In nondeterministic environments, the exact state of the world cannot be determined, which makes utility functions rely on making best-guess decisions.

- Episodes

Episodic environments are those where the agent's current action choice does not rely on its past actions. Non-episodic environments require agents to plan for the effects of its current actions.

- Type of environment

Static environments are where the environment doesn't change while an agent is making a decision. Dynamic environments can change during processing, so agents may need to feed back into the environment or anticipate changes between the time of input and output.

- Flow of data to environment

One critical aspect of agent design is how the data is received. In a chess game, for example, there are only a finite number of moves that the agent can select. This would be a discrete environment, as there are only so many moves that can be performed. On the other hand, on the operating table, for example, things can change at the last millisecond. Continuous data environments are environments in which data seems infinite or without end.

An intelligent agent is fundamentally one that can display characteristics such as learning, independence, adaptability, inference, and to a degree, ingenuity.

For this to occur, an agent must be able to determine sequences of actions autonomously through searching, planning, adapting to changes in its environment, and from feedback through its past experiences. Within machine learning, there are further terminologies.

Training Data

Training data is the data that will be used by the learning algorithm to learn possible hypotheses. An example is a sample of x including its output from the target function.

Target Function

This is the mapping function f from x to $f(x)$.

Hypothesis

This is an approximation of f .

Learner

The learner is the learning algorithm or process that creates the classifier.

Hypothesis

A hypothesis is the possible approximation of f that the learning algorithm can create.

Validation

Validation includes methods used within machine learning development that provide a method of evaluation of model performance.

Dataset

A dataset is a collection of examples.

Feature

A feature is a data attribute and its value. As an example, skin color is brown is a feature where *skin color* is the attribute and *brown* is the value.

Feature Selection

This is the process of choosing the features required to explain the outputs of a statistical model while excluding irrelevant features.

What Is Machine Learning?

In 1959, Checkers' creator Arthur Samuel defined machine learning as a “field of study that gives computers the ability to learn without being explicitly programmed”.[47]

Machine learning was born from pattern recognition and the theory that computers can learn without being programmed to perform specific tasks—that is, systems that learn without being explicitly programmed. As a result, learning is driven by data—with intelligence acquired through the ability to make effective decisions based on the nature of the learning signal or feedback. The utility of these decisions is evaluated against the goal.

Machine learning focuses on the development of algorithms that adapt to the presentation of new data and discovery. Machine learning exemplifies principles of data mining but is also able to infer correlations and learn from them to apply to new algorithms. The goal is to mimic the ability to learn in a human, through experience, and achieving the assigned task without, or with minimal, external (human) assistance.

Just as with learning in humans, machine learning is composed of many approaches. At its most basic, things are memorized. Second, we learn from extracting information (through reading, listening, learning new things). And third, we learn from example.

For example, if we were being taught square numbers; and if students were shown a set of numbers $Y = \{1, 4, 9, 16, \dots\}$ and taught that $Y = n^2$; they would typically understand the notion of square numbers without being explicitly taught every square number.

How Is Machine Learning Different from Traditional Software Engineering?

Traditional software engineering and machine learning share a common objective of solving a problem or set of problems. The approach to problem-solving is what distinguishes the two paradigms.

Traditional software engineering or programming refers to the task of computerizing or automating a task such that a function or program, given an input, provides an output. In other words, writing a function f , such that given input x , output $y = f(x)$. This is done with logic, typically if-else statements, while loops and Boolean operators.

Machine learning differs from conventional programming in that rather than providing instructions about the function f , the computer is provided input x and output y and expected to determine or predict function f (see Figure 3-2). Whereas traditional programs have been written by humans to solve a problem, machine learning programs learn through reasoning to solve a problem from examples, rules, and information. Many classifiers in machine learning assist with prediction. Machine learning programs can also learn to generalize and help with issues of uncertainty due to the use of statistics and probability-driven techniques. Models can learn from previous computations or experience to produce reliable, repeatable decisions and results.

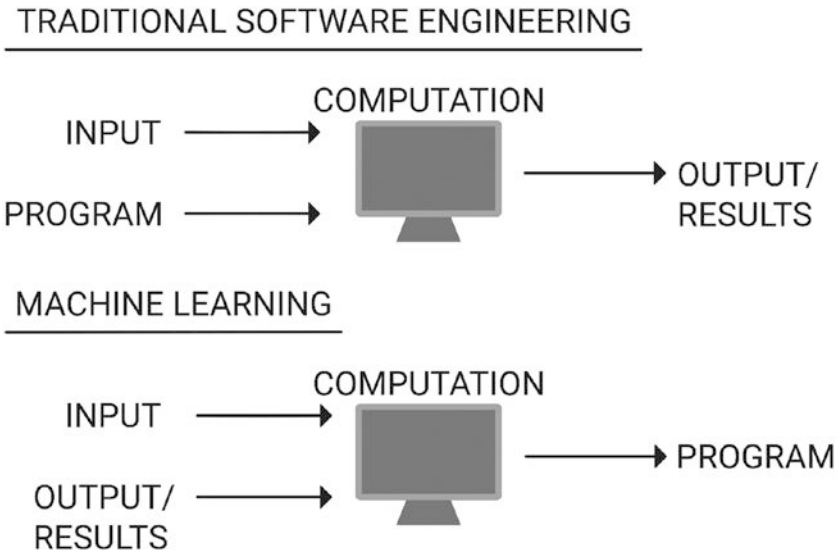


Figure 3-2. *How machine learning works*

The availability of data, information on the Internet and leaps in the amount of digital data being produced, stored, and made accessible for analysis is enabling engineers to realize machines that can learn what we prefer and personalize to our requirements. Machine learning can be seen in web search rankings, drug design, personalization of preferences and social network feeds, web page design, self-driving cars, and virtual assistants.

Machine Learning Basics

Tasks for machine learning algorithms are broadly categorized as the following (see Figure 3-3):

- Supervised learning (also known as inductive learning)
- Unsupervised learning
- Semi-supervised learning
- Reinforcement learning

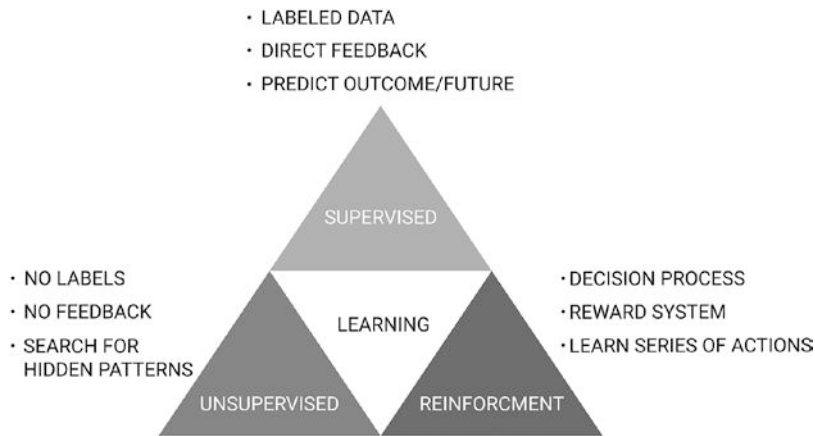


Figure 3-3. Methods of learning

Supervised Learning

In supervised learning, algorithms are presented with example inputs and the subsequently desired outputs in the form of training data, with the goal to learn general rules that map inputs to outputs (see Figure 3-4). Input data is called training data and has a known output (result) associated with it. The training data guides the development of the algorithm. A model is created through a training process in which the model makes predictions and is corrected when predictions are incorrect.

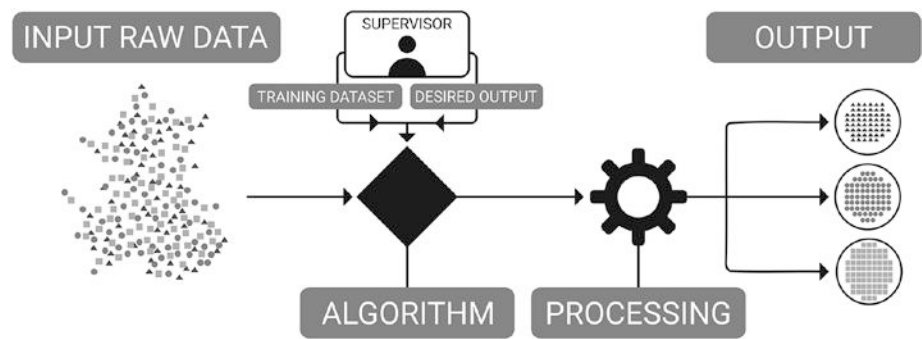


Figure 3-4. Supervised learning: how it works

Training continues until the model achieves the desired level of accuracy on the training data. Supervised algorithms include logistic regression. Supervised machine learning algorithms can apply what has been learned in the past to new data using labeled data examples to predict future events. The learning algorithm can also compare the output with the correct output to find any errors and modify the model accordingly.

Supervised learning problems can be employed in the following forms:

- **Classification:** To predict the outcome based on a training dataset where the output variable is in the form of distinct categories. Models are built through inputting training data in the form of prelabeled data. Classification techniques define decision boundaries and include support vector machines, naïve Bayes, Gaussian Bayes, k-nearest neighbors (KNN), and logistic regression. An example of classification is the label (or in real-world application, diagnosis) of someone as sick or unhealthy based on a set of symptoms.
- **Regression:** Very similar to classification. The only difference between classification and regression is that regression is the outcome of a given sample where the output variable is in the form of real values (i.e., the temperature as a value, rather than a classification of hot or cold). Examples include height, body temperature, and weight. Linear regression, polynomial regression, support vector machine (SVR), ensembles, decision trees, and neural networks are examples of regression models.
- **Forecasting:** The method of making predictions based on past and present data. This is also known as time-series forecasting.

- An ensemble is a type of supervised learning that combines multiple, different, machine learning models to predict an outcome on a new sample. In an ensemble, models become features.

Unsupervised Learning

When people reference systems that can learn for themselves, they are actually referencing unsupervised learning (see Figure 3-5). In unsupervised learning, the learning algorithm doesn't receive labels for data, leaving the algorithm to find structure from the input. As data is unlabeled, there is no evaluation of the accuracy of the structure that is output by the algorithm. Data may be missing both classifications and labels. As a result, the model is developed through interpretation: through finding hidden structures and inferences in the input data. This may be through the extraction of rules, reducing data redundancy, or as the result of organizing data.

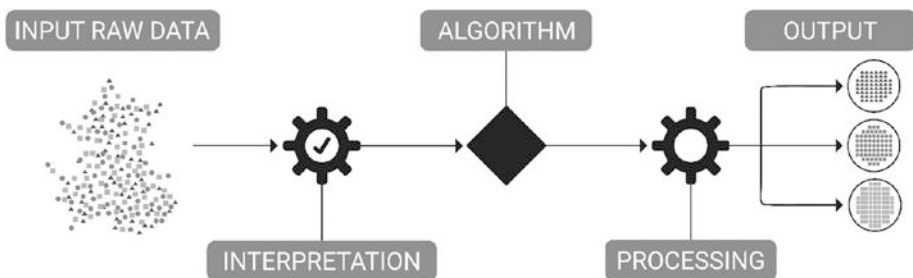


Figure 3-5. *Unsupervised learning: how it works*

This includes clustering, dimensionality reduction, and association rule learning. The algorithm may never find the right output but instead models the underlying structure of the data.

There are three types of unsupervised learning problems:

- **Association:** Discovering the probability of the co-occurrence of items in a collection. This is used extensively in marketing as well as healthcare. For instance, this would be the percentage likelihood of developing any form of cancer if diagnosed with poorly controlled type 2 diabetes. Association is similar to classification; however, any attribute can be predicted in association, whereas classification is binary.
- **Clustering:** Grouping items such that items within the same cluster are more like each other than to items from another cluster.
- **Dimensionality reduction:** Dimensionality reduction can be achieved through feature selection and feature extraction. Dimensionality reduction may mathematically re-represent, or transform data.

Feature extraction performs data transformation from a high-dimensional to a low-dimensional space. This can involve reducing the variables of a dataset while maintaining data integrity and ensuring that the most important information is represented.

Feature extraction methods can be used for dimensionality reduction. Typically, a new set of features would be created from the original feature set. An example would be combining all of a patient's clinical test results into a health risk scoring from the results that are demonstrated to affect mortality. Reducing the number of dimensions allows easier visualization, particularly in two or three dimensions, as well as reducing time and space required for storage.

Algorithms include Apriori algorithm, FP-growth (also known as frequent pattern growth), Hidden Markov Model, Principal component analysis (PCA), singular value decomposition (SVD), k-Means, neural networks, and deep learning. Deep learning utilizes deep neural network architectures, which are types of machine learning algorithms.

Semi-supervised

Semi-supervised learning is a hybrid where there is a mixture of labeled and unlabeled data in the input. Although there may be a desired outcome, the model also learns structures to organize data and make predictions. Example problems are classification and regression.

Reinforcement Learning

This is where systems interact with a dynamic environment in which an agent must perform a specific goal. Reinforcement learning acts as an intersection between machine learning, behavioral psychology, ethics, and information theory.

The algorithm is provided feedback concerning rewards and punishments as it navigates the problem. Reinforcement learning allows the agent to decide the next best action based on its current state and by learning behaviors that will maximize the reward. Optimal actions (or an optimal policy) are typically learned through trial and error and feedback. This allows the algorithm to determine ideal behaviors within context.

Reinforcement learning is typically used in robotics: for instance, a robotic hoover that learns to avoid collisions by receiving negative feedback through bumping into tables, chairs, and so forth—although today, this also involves computer vision techniques. Reinforcement learning differs from standard supervised learning in that correct examples are never presented, nor are suboptimal decisions explicitly corrected. The focus is on performance in the real world.

Similar to how humans learn, reinforcement learning assumes rational behavior regarding learning from one's experience, and as such is assisting in understanding how humans learn from experience to make good decisions (see Figure 3-6).

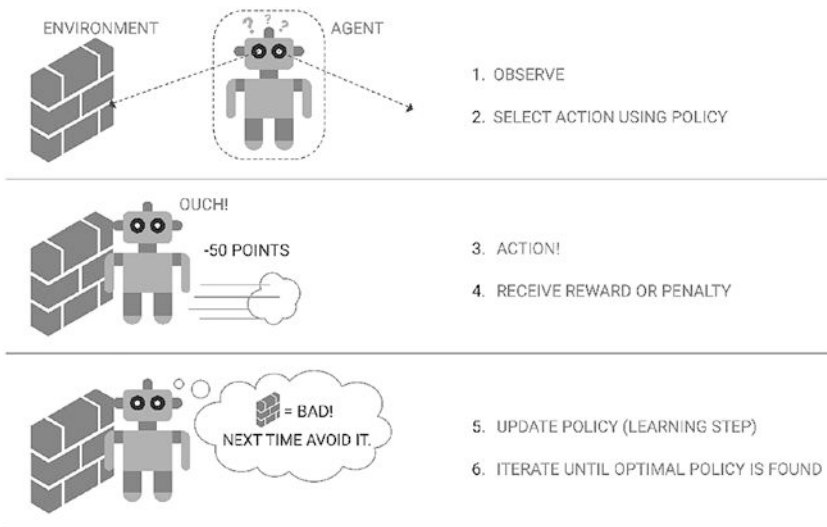


Figure 3-6. Reinforcement learning: how it works

DeepMind, for example, exposed their AlphaGo agent to master the ancient Chinese game of Go through showing “AlphaGo a large number of strong amateur games to help it develop its own understanding of what reasonable human play looks like. Then we had it play against different versions of itself thousands of times, each time learning from its mistakes and incrementally improving until it became immensely strong.” [48]

Reinforcement learning is used in autonomous vehicles. In the real world, the agent requires consideration of aspects of the environment including current speed, road obstacles, surrounding traffic, road information, and operator controls. The agent learns to act in a way (or develop a policy) that is appropriate for the specific state of the

environment. This leads to new ethical dilemmas posed by AI and machine learning for agents with reward-driven programming.

Data Mining

Data mining, also known as knowledge discovery in databases, is defined as the process of exploring data to extract novel information (including patterns and relationships) using sophisticated algorithms. This can be applied to every aspect of life and industry and enables the following:

- Pattern predictions based on trends and behaviors
- Prediction based on probable outcomes
- Analysis of large datasets (particularly unstructured)
- Clustering through identification of facts previously unknown

Data mining uses the power of machine learning, statistics, and database techniques, and is at the heart of analytics efforts (see Figure 3-7). Data mining typically uses large amounts of data and can also give simpler, descriptive analytics. For instance, in healthcare, data mining applications can be used to identify conditions or medical procedures that are often seen together. Data mining also allows forecasting, through identifying relationships within the data. This can help insurance companies detect risky customer behavior patterns.

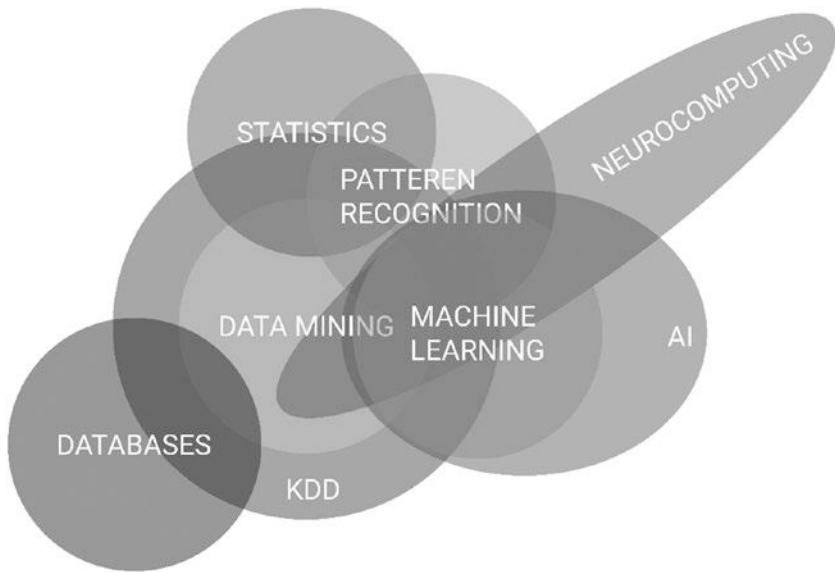


Figure 3-7. *Data mining and its wider synergies*

Machine learning can be used for data mining. Data mining will inform decisions such as the kind of learning models best suited for solving the problem at hand. However, data mining also encompasses techniques besides those encapsulated within machine learning. In data mining, the task is of unknown knowledge discovery, whereas machine learning performance is evaluated with respect to reproducing known knowledge.

For instance, if you had a dataset of a patient's blood pressures, you can perform anomaly detection, which is considered a data mining task to identify previously unknown outliers. The task may make use of machine learning techniques, perhaps a k-means algorithm for cluster analysis, to identify the anomalous data and assist the algorithm's learning.

Data mining is not the same as natural language processing (NLP). NLP may be a technique used in data mining to help an agent understand or "read" text. Data mining algorithms include C4.5, k-means, Support vector machines, Apriori, EM (expectation-maximization), PageRank, AdaBoost, kNN, Naive Bayes and CART (Classification and Regression Trees).

Parametric and Nonparametric Algorithms

Algorithms may be parametric or nonparametric in form.

Algorithms that can be simplified to a known, finite form are labeled parametric, whereas nonparametric algorithms learn functional form from the training data. In other words, with nonparametric algorithms, the complexity of the model grows with the amount of training data. With parametric models, there is a fixed structure or set of parameters. As a result, parametric models are faster than nonparametric models. However, nonparametric models can provide better accuracy if they are supplied enough training resources in terms of training data and time (see Figure 3-8).

	 PROS	 CONS
PARAMETRIC ALGORITHMS	<p>SIMPLER</p> <p>SIMPLER TO UNDERSTAND AND INTERPRET</p> <p>FASTER</p> <p>QUICK TO FIT TO DATA</p> <p>LESS DATA</p> <p>REQUIRES LESS DATA TO PERFORM WELL</p>	<p>LIMITED COMPLEXITY</p> <p>SUITED TO SIMPLER PROBLEMS WHERE STRUCTURE IN THE DATA CAN BE INFERRED</p>
NONPARAMETRIC ALGORITHMS	<p>FLEXIBILITY</p> <p>CAN FIT VARIETY OF FUNCTIONAL FORMS, WHICH DO NOT NEED TO BE ASSUMED</p> <p>PERFORMANCE</p> <p>PERFORMANCE WILL LIKELY BE HIGHER THAN PARAMETRIC ALGORITHMS AS SOON AS DATA STRUCTURES GET COMPLEX</p>	<p>SLOWER</p> <p>COMPUTATIONS ARE SIGNIFICANTLY LONGER</p> <p>MORE DATA</p> <p>REQUIRES SIZEABLE DATA TO LEARN</p> <p>OVERFITTING</p> <p>AFFECTS MODEL PERFORMANCE</p>

Figure 3-8. *Parametric vs. nonparametric algorithms*

Linear models such as linear regression and logistic regression are parametric, whereas K-nearest neighbor, neural networks, and decision trees are nonparametric.

How Machine Learning Algorithms Work

Supervised learning is the computerized task of inferring a function from labeled training data.

$$Y = f(x),$$

which can be understood as

$$\text{Output} = \text{function}(\text{Input}).$$

Inputs and outputs can be referred to as variables and are typically vector in format: f refers to the function, which you are trying to infer from the Input.

Take the example of a patient who is to be prescribed a particular treatment. In real life, the process may include the patient looking at research describing qualities about the medication or treatment he or she was recommended. For example, if he or she were to see that the research mostly consists of words like “good,” “great,” “safe,” and so forth, then we would conclude from the sentiment that the treatment is a good treatment, and we can feel confident in accepting it. Whereas if words like “bad,” “not good quality,” and “unsafe” were to appear regularly, then we conclude that it is probably better to look for another treatment. The research helps us perform an action based on the pattern of words that exist in the treatment research.

Machine learning attempts to understand this human decision-making process through algorithms. It focuses on the development of algorithms (or programs) that can access data and use it to learn for themselves. Machine learning is very good for classification and prediction tasks. Machine learning is most commonly used for predictive analytics—learning the mapping of $Y = f(X)$ to make predictions of Y for new X .

Practical examples of machine learning are the following:

- Fraud detection x is the property of the transaction $f(x)$ is whether a transaction is fraudulent
- Disease diagnosis x is the properties of the patient $f(x)$ is whether a patient has disease
- Speech recognition X are the properties of a speech input $f(x)$ is the response given to the instruction contained within the speech input

The process of machine learning begins with input data provided as examples, direct experience, or instructions—to identify patterns within the data and make better decisions in the future based on the data that was provided. The aim is to enable the program to learn automatically without human intervention or assistance and adjust rational actions accordingly.

How to Perform Machine Learning

Performing machine learning is like typical algorithm methodologies. It involves the following:

1. Specifying the problem as a learning task(s)
2. Preparing the data
3. Choosing the learning method
4. Applying the learning method
5. Assessing the method and results
6. Optimization
7. Reporting the results

Specifying the Problem

The first thing when it comes to machine learning is specifying the problem. This involves understanding what the problem is, how we can solve the problem, and how it can be evaluated. It is useful to understand why you want to solve the problem, especially if you are creating a budget case for a machine learning project. This includes evaluating the following:

- Why is it important?
- What do you hope to achieve? What are your best hopes?
- What are the inputs and outputs for the task at hand?
- Is this data available?
- How will the results be beneficial?
- Is this exploratory?
- What are the KPIs (key performance indicators); how is performance measured?
- What does success look like?
- Do we have the talent on hand to attempt the task?
- What are the limitations to the task at hand? This could include time, finance, skillsets, experience, domain knowledge, data availability, and so forth.

To solve the problem, data available to you will include examples and background information.

Examples

Categorization is a skill that can be learned through examples. It can be used for predictive tasks such as disease risk or resource demand prediction. Certain problems present themselves in traditional healthcare settings, such as data from medical appointments and certain diagnoses having complementary handwritten notes. Bad examples can help to introduce error. For classification, examples are typically provided in the form of positives and negatives.

Background Information

This refers to the knowledge and axioms associated with a machine learning problem. This may include metadata, attributes, or relationships between concepts. For example, in a hypertension prediction tool, the diagnosis of type 2 diabetes would be related to the likelihood of hypertension based on medical evidence and data. This is information that is utilized in the concepts learned to achieve categorizations.

Errors in the Data

Be mindful of the fact that errors from the real world can creep in: for example, from incorrect classifications, missing data, incorrect background information, and repeated data. Errors could creep in, for example, from the digitalization of paper documents and the machine-reading of written text. It is important to remember the model is only as good as the data it is provided. It is worth spending time verifying the data to ensure that human errors are minimized.

Preparing the Data

Machine learning algorithms learn from the data they are trained on. It is mission critical to provide the model with valid, correct data to learn. Data must be prepared in a usable format. In a real-world scenario, this would typically involve understanding the data that would be used to model the problem and exporting the data. The data must then be processed to ensure correct formatting, removal of erroneous data, and the fixing of any missing data. The dataset size may be more than required, and so dataset sampling may also be required.

Every machine learning candidate will acquire skills in data preparation. Data preprocessing is essential to have tidy, valid data. Tidy, valid data is key to having robust, veracious outcomes.

First, we need to select the data. When first extracted, the data is raw. You may need to tidy your data. You may need to clean your data, which involves removing or fixing bad data.

Attribute Selection

Attribute selection is also called variable selection or feature selection. It is essentially filtering—and refers to the selection of a subset of the original example set that is most relevant in the predictive modeling at hand. Feature selection is different from dimensionality reduction in that dimensionality reduction methods create new combinations of attributes. Feature selection includes and excludes attributes rather than creating new ones.

This is where domain knowledge is advantageous. Understanding the context in which the problem fits enables you to better determine a feature set.

Transforming Data

In most industries, but particularly in the healthcare setting, many instances require anonymized data, and thus some features may be completely removed from the dataset. Formatting may need to be changed; for example, you may need to export your data to something that is a flat file rather than a traditional relational database. You may only want to work with a subset of your sample data.

Digitalized text, such as handwriting, may have errors. Check your datasets for errors, biases, and inconsistencies. This is a very common problem.

Data may also need to be transformed. This is typically guided by the algorithm you are using and the data available. For instance, if a machine learning algorithm was taking the input of a diagnosis of type 2 diabetes as an integer “[0 for false]” and “[1 for diagnosed with type 2 diabetes],” but the medical notes read specified “type 2 diabetes” or “no type 2 diabetes,” the data would need to transform from text to integers to be machine readable. Scaling may also be used—for example, converting metric to imperial units.

Transformations can include the following.

Aggregation

This refers to attributes that can be aggregated into an individual attribute. For example, a log of particular patient attendance to health education could be aggregated into a count of the total number of attendances rather than representing each attendance separately.

Decomposition

This refers to partitioning attributes such that they may become more useful to a machine learning technique. For example, a patient’s records may have their ethnicity listed as Asian. However, it may be better to know

the country of birth and primary language of the patient to more accurately predict criteria such as type 2 diabetes risk.

Scaling

Data can contain attributes with varying quantities. For example, weight in kilograms, height in centimeters, and clinical markers may use different scales completely.

Tidy data is the result of data preprocessing including spot checks, formatting, and transformations.

Machine learning can only be as good as the data you use to train it. If the data is biased (i.e., not representative), the learning will be biased. At the same time as you are exploring datasets, be careful to avoid introducing human error. This is typically satisfied through ensuring good version control.

Weightings

Weight refers to the priority or influence of a particular feature. Typically, positive weightings would increase the prediction value, whereas negative weights would reduce the prediction value.

How Much Data Do I Need?

Sadly, there is no definitive answer to this question—it depends. This is something that you explore through training your model and deploying it in the real world. Several factors can affect the “learning curve.”

These include the complexity of the problem and the complexity of the algorithm. This, in turn, affects the amount of data required. Whether the model can make outputs that have veracity in the real world will depend on the sophistication of the algorithm and quality of data provided.

Assess the availability and usefulness of more data. In a healthcare setting, randomized controlled trials on medications rarely assess more than 1,000 patients in any given setting. In a surgery setting, access to bigger data is likely but may not be required for training a model.

Nonlinear, nonparametric algorithms require more data to become better as to their accuracy. These are typically the more powerful machine learning methods. However, due to the validity of the data affecting both bias and variance, it's key to train and validate the model robustly on as much data as efficiently possible.

Choosing the Learning Method

The standard question any newcomer to machine learning has is “which algorithm should I use?” The algorithm of choice depends on several factors including the size, quality, and nature of data; the task deadline; and available resources and motivation for using the data. The learning technique is also referred to as the “representation of the solution,” as each machine learning approach represents data differently.

It is almost impossible to predict which approach will best perform on the data. That isn't to say that it will perform well in the real world! Still, there are some general algorithms that can be used for certain problems, which are highlighted in [Chapter 4](#).

NO FREE LUNCH

The “No Free Lunch” theorem in machine learning states that there is no one algorithm that works best for every problem. This is particularly relevant for supervised learning, as there are many factors that affect the performance of an algorithm.

Applying the Learning Methods

Machine learning tasks are typically conducted in a variety of programming languages: predominantly R, Python, Matlab, and SQL. Java and C are also commonly used.

- R is typically used for statistical analysis. It allows you to understand and explore data using statistical methods and graphs and contains an extensive range of machine learning algorithms.
- Python is a language well suited to machine learning. Extensions such as NumPy and SciPy are particularly useful for machine learning and data analysis.
- Matlab is the language many university students begin with. It is useful for fast prototyping, as it contains a large machine learning repository.
- SQL is a language used for managing data held in a traditional database management system.

Should I Code My Machine Learning Algorithm from Scratch?

Several decisions need to be made from a resource and value perspective that affect the approach. A proof of concept project may be short of time, whereas a highly specific and complex task will contain a blend of approaches.

It is entirely possible to run a machine learning algorithm through implementing the ready-made algorithms provided by learning libraries such as scikit-learn, SciPy, Pandas, Matplotlib, TensorFlow, Keras, and so forth.

This has been the traditional approach to machine learning: establishing a data science team that includes members experienced in languages such as Python and R. Members with experience in machine learning will typically be more impactful than those who are beginners to machine learning concepts. Engineers develop an understanding of the appropriate machine learning techniques to employ given certain problems through experience. This experience can prove vital in avoiding wasted time, resources, and exploration. Engineers with fundamentals in computer science will typically be familiar with many of the concepts present in machine learning. Choosing a programming language can influence the APIs and standard libraries you can use in your implementation. Enthusiasts may want to avoid using ready-made algorithms and coding the implementations themselves to gain greater knowledge of the mathematics, statistics, and logic involved.

In many instances where time is constrained, time for engineering enthusiasm is limited. As a result, cloud-based services such as Google AI, Microsoft Azure, IBM Watson, and a host of other providers allows the provider to take care of the algorithm; and many provide visual interfaces that enable the user to attempt particular machine learning techniques on the data uploaded. Cloud-based vendors offering data mining and machine learning applications are growing and available at a low cost. One of the benefits of using such vendors is the ability to try more than one type of model for an experiment and compare the results, allowing you to find the best solution for your problem task.

For those with more traditional programming experience, almost all offer API (or web service) creation facilities, allowing you to create an interface to engage with the hosted model. All that is required to integrate with APIs is an efficient traditional programmer or web developer.

As the world adopts open source, consider browsing for implementations that already exist through GitHub, Reddit, and other machine learning forums.

API stands for application programming interface. APIs enable communication between software components, and in particular, can be used to interface between cloud-based machine learning providers and in-house development teams. There are many public APIs that can be used to assist in your machine learning tasks.

Embedding machine learning APIs should be approached with diligence and discretion. In an increasingly open source environment, as well as good documentation and functionality, ensure there are visible displays of popularity and robustness such as GitHub status, search engine popularity, and conversations.

Training and Test Data

A test set and training set is selected from the prepared data. The algorithm is trained on the training dataset and evaluated against the test dataset.

In many cases, it is a case of trialing several machine learning methods. It is useful to understand two terms when it comes to machine learning modeling.

- Signal: the true underlying pattern in a dataset
- Noise: random or irrelevant patterns in a dataset

Some machine learning techniques may return one solution, whereas others may produce several. It is often a case of gathering their outputs (also referred to as hypotheses, or learned models) and evaluating their outputs.

The assessment of hypotheses is conducted through evaluating the predictive accuracy, comprehensibility, or utility. In most cases, Occam's razor is used, where the simplest solution is chosen if all else is equal.

- Predictive accuracy refers to the accuracy at which the agent performs the task of classification.

- Comprehensibility refers to how well we as humans can understand the output. For example, two Facebook bots recently began talking to each other in a language of their own.[49]
- Utility refers to the problem-specific measure of worth. For example, in drug synthesis, a combination of compounds that are unsafe for humans may be the best combination. However, this may not have the greatest accuracy. Equally, this may override accuracy and comprehensibility.

After the assessment is complete, a candidate hypothesis is chosen.

There are various ways to partition data into training and test datasets to train the machine learning model.

Hold-Back Method

The hold-back method is a simple training method that can be employed that involves holding a proportion of the data to test on, leaving the remaining data for training and validating the model. Testing will typically give an indicator of the model's effectiveness or accuracy.

n-fold Cross-Validation

The n-fold cross-validation involves splitting the dataset into equally sized groups of instances (or folds). The model is then trained on all but one fold and tested on the final, left-out fold. This process is repeated, and each fold is left out for one iteration.

The number of folds can vary based on the size of the dataset. Common folds include three, five, seven, and ten folds. The aim is to balance between the size and representation of data in the training and test datasets. It enables the model to be tested n times on different datasets. This enables us to make the most out of the data, training n times on it.

Some machine learning texts may refer to cross-validation as the cross-validation stage of the machine learning process. This is because machine learning engineers often use this approach to reduce the likelihood of overfitting.

Monte-Carlo Cross-Validation

Monte-Carlo cross-validation is similar to n-fold cross-validation. It involves randomly splitting the dataset into training and test data. The model is tuned to the training data, and predictive accuracy is evaluated using the test data. Results are then averaged over the splits.

Try Many Algorithms

Try a number of techniques that are relevant to your problem. This allows you to compare and choose the solution judged on accuracy, comprehensibility, and utility. Machine learning software typically supports multiple algorithms. Find the best algorithm for the job by trying as many as feasible.

Assessing the Method and Results

The performance of machine learning tasks varies on the representation of data given. For example, a patient record analyzed by an AI system does not examine the patient directly. Instead, data is fed into a system, with each piece of information relating to the representation of the patient known as a feature. It is not necessary to require complete feature sets as part of representations to have highly confident outputs.

The goal of a machine learning algorithm is to generalize well, neither underfitting nor overfitting. Generalization refers to the model performing with maximum accuracy on instances not seen during training.

Algorithm Accuracy Evaluation

Many supervised learning problems have binary classifications, where the problem is to learn a way of classifying unseen data into one of two categories. These are known as positive and negative categories. A false positive is received if the output from a binary classification agent that is given new data should have been categorized as negative but was classified as positive. A false negative is when an agent categorizes new data as negative but is incorrect.

In many medical settings, a false positive is not as bad as having a false negative. For example, false positives in this instance could include the diagnosis of illness for a patient based on a set of particular symptoms. It is evidently better to be told that you have a condition and not have it in reality, rather than being told you do not have a condition, only to be told that in fact, you do. A false positive would mean that someone was not diagnosed, which is perhaps more concerning.

A simple way of measuring predictive accuracy of a particular hypothesis over a test dataset is to calculate the number of correctly identified outcomes (both positive and negative). For example, if a heart disease tool learned a hypothesis and was given 250 new candidates to classify as having heart disease, and correctly classified 133 of 150 positives and 98 of 100 negatives, it would have an accuracy of

$$(133 + 98)/250 = 92.4\%$$

This would be a 92 out of 100 chance of correctly categorizing an unseen example.

Generalization is measured by accuracy. Both under- and overfitting cause poor model performance and are detrimental to accuracy. If a model that was classifying with 92% accuracy in testing but only 30% on unseen data, we would say that the model doesn't generalize well from training to unseen data.

This is overfitting. Overfitting is when a model fits the training data too well. It memorizes answers rather than learning concepts from them. In this case, the model cannot distinguish signal from noise and learns the noise present in the data. As a result, the model does not understand the actual patterns, or signals, in the data. As such, the model does not apply to new data and is determinantal to the model's ability to generalize. Decision trees are examples of algorithms that can overfit training data. This can be improved through pruning. In many cases, it's a case of keeping things simple.

If a model performs better on training datasets than on unseen test sets, then the model is likely to be overfitting. Cross-validation is a technique to employ to measure against overfitting. Further still, training with more data may help algorithms infer hypotheses better.

The concept of "more data gives better results" is misleading. In many cases, more data may help. However, more data does not imply good data. If data with noise were to be added, it wouldn't help reduce overfitting. Hence, this is why it's paramount to ensure data is clean and appropriate.

Data leakage occurs when data you are using to train your machine learning model has the information you are trying to predict in it. This leads to poor generalization and contributes to overfitting. Data leakage can occur from test data leaking into training data or additional features. A trivial example of data leakage would be a model that uses the response variable itself as a predictor, thus concluding, for example, that "a patient with type 2 diabetes has type 2 diabetes."

If the performance of your machine learning algorithm is too good to be true, data leakage may be the reason it is. An n-fold cross-validation can help to reduce data leakage.

Underfitting is simple to detect, as it will have poor performance. It refers to a model that cannot model the training data or generalize to new data. Typically, this would be represented by simple models that have too few features. Results would tend to show bias over variance (i.e., wrong outcomes).

Bias and Variance

Supervised machine learning aims to understand the signal from the dataset while ignoring the noise. There is a machine learning dilemma of minimizing the two sources of prediction error that prevent supervised learning algorithms from generalizing beyond their training set. There is a trade-off between a model's ability to minimize bias and variance (also known as the bias-variance trade-off).

Bias

Bias refers to the error the model has from learning from erroneous data. Although this sounds strange, datasets may present randomly. Bias measures how far off, in general, these models' predictions are from the actual correct value, or true signal.

- Low bias: Better understanding of true signal from a dataset
- High bias: Worse understanding of true signal from a dataset

High bias can often indicate that a better machine learning technique is available.

Variance

This refers to error from sensitivity to fluctuations in the training set.

- Low variance algorithms output similar models based on the dataset.
- High variance algorithms output sufficiently different models based on the dataset.

Bias and variance can best be illustrated using a dartboard analogy (see Figure 3-9). Take a dartboard where the bull's-eye of the dartboard represents a model that perfectly predicts correct values. As you move further away from the bullseye, the predictions become worse. If each new model was represented by a new dart, each throw represents a realization of the model given the variability in the training data. With a good distribution of data for training, models can predict well. Data with anomalies and redundancies will result in poor predictions. This can be visualized on a dartboard.

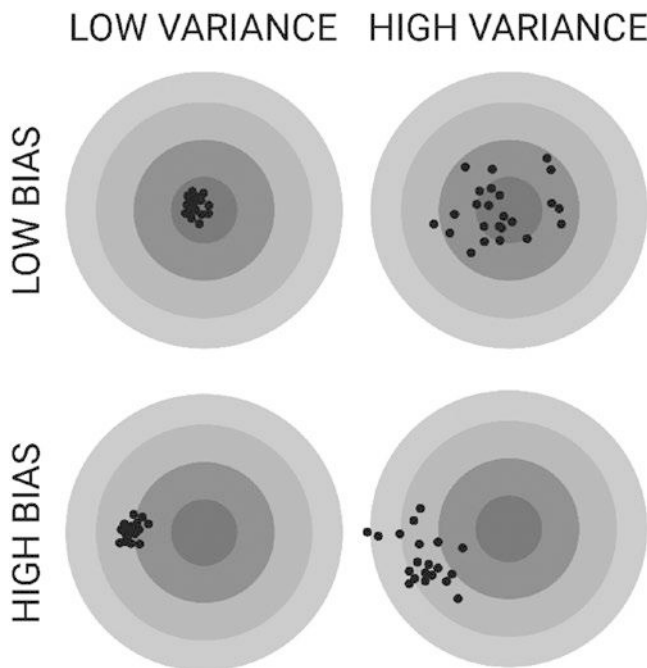


Figure 3-9. *Bias and variance illustrated*

The trade-off engineers must make with machine learning algorithms is that an algorithm with low bias must be flexible to enable it to fit the data well. However, if there is too much flexibility, each dataset will be

interpreted differently and have high variance. There are several ways to reduce algorithm complexity, which are discussed in Chapter 4.

Many supervised learning toolkits provide methods to control the bias-variance tradeoff either automatically or by providing a parameter that can be adjusted.

Low-variance algorithms tend to have reduced complexity when compared to low-bias algorithms. Low-variance algorithms have a simple or rigid structure that results in models that are consistent but inaccurate on average. This includes algorithms like regression, naïve Bayes, and so forth.

Low-bias algorithms tend to be more complex with a more flexible structure. As a result, they are accurate on average; however, they can be inconsistent. This includes nonparametric algorithms such as decision trees and k-nearest neighbor.

The key in your project is to find a balance between bias and variance that minimizes total error (see Figure 3-10).

$$\text{Total Error} = \text{Bias}^2 + \text{Variance} + \text{Irreducible Error}$$




Figure 3-10. Finding optimal algorithm balance

The irreducible error refers to error regardless of the model or dataset that is used. This is the noise that comes from the training data. It is a constant, as real-world data always has some degree of noise.

Performance Measures

Examples of methods for evaluating algorithms include accuracy, prediction and recall, squared error, likelihood, posterior probability, cost, and entropy k-L divergence.

Optimization

Just because a model performs well doesn't mean it is the best and only model. The key is to get the maximum veracity from your results. Techniques such as cross-validation are useful to determine confidence in model results. However, we can optimize machine learning algorithms further.

Besides, we may want to improve the performance regarding time to output.

Algorithms can be optimized through the following:

- Algorithm tuning

Tuning can be understood as the process of optimizing the parameters that impact the model to enable the algorithm to perform the “best,” based on what “best” is specified as.

- Training and validation data

Approaching a machine learning problem with an n-fold cross-validation technique allows you to validate your results; whereas a hold-back methodology has a waterfall approach to learning.

- Evaluating a range of methods

Trialing a variety of machine learning methods is useful to determine which can achieve the most accurate results.

Improving Results with Better Data

There are some ways to improve performance with data, including the following:

- Getting more data

If you can get more data, this may help the model to improve its performance.

- Getting better quality data

Rather than more data, it's often better, where possible, to get better quality data. Typically, this would improve the signal within the data and reduce the noise. You can also enhance the quality of data through cleaning the data.

- Resample data

Partitioning the sample data into different sizes or distributions may represent the concept better or help improve performance by reducing attributes (features).

- Data representation

Data can be represented in different ways through the varying machine learning methods. Changing the type of machine learning method applied and reevaluating the problem may optimize your model.

- Feature selection

Addressing features that are and are not important can present new inferences for your model to learn. Increasing the number of features may not always be beneficial. However, in unsupervised learning instances, it may be beneficial to include all features to ensure nothing goes unexplored.

- Feature engineering

The key to success is the representation of data. Feature engineering involves exposing the data in other ways through decomposition and aggregation. For example, aggregation may involve categorizing patient ages in brackets of 10 years, rather than specifying a continuous number. Decomposition involves reducing dependencies within the data. This may even involve including attributes previously excluded and unseen to the model during data preparation. One of the most famous quotes behind the power of data is that of Google's Research Director and machine learning pioneer Peter Norvig, claiming "We don't have better algorithms. We just have more data."

Should I Choose a Supervised or Unsupervised Algorithm?

When developing models, it's rarely a case of supervised or unsupervised learning. In fact, they can often be used together. Unsupervised learning is a useful technique for dimensionality reduction and feature engineering, which can be used in supervised learning problems.

Ensembles

Ensembles are created when several classifiers specialize in various aspects of the problem. Ensembles turn a model into a feature. There are three types of ensemble algorithms known as bagging, boosting, and stacking, covered in Chapter 4.

Ensembles are for advanced practitioners of machine learning.

Problem Distribution

Another optimization method is distributing the processing used for a machine learning task. Tools from Microsoft, such as Azure, and from Google, such as AI.io, are examples of warehouses that can distribute

problems at scale. In reality, most problems can be contained on a multicore implementation.

Implementation Problems

The user interface or UI is the key component in a machine learning application that has human interaction. This is not only where results will be displayed but where the most intelligence can be gathered. Rethinking the user interface and user experience is worth considering when engaging with people.

Reporting the Results

In many settings, this is the most crucial aspect to ensure internal stakeholders support machine learning projects. Typically, a PowerPoint presentation and white paper will be good enough to present to your company or clients.

Frame the problem as a scientific experiment, assessing the following:

- Why—define the context of the problem and the motivation
- Problem—describe the issue as a question
- Solution—explain the answer to the questions proposed
- Findings—describe the discoveries made in the data and inferences made by the models
- Conclusions—evaluate what was learned from the project, addressing the limitations, ethical concerns, privacy, and data sensitives (particularly with GDPR).

Once you've reported your findings, it's up to your stakeholders to determine the value of the project before deployment, if deployed at all.

Topics for discussion include how to handle the reporting of false negatives, false positives, and the accuracy of predictive claims. Also, the implications of predictions are a crucial consideration. Some concerns may be first-time problems. For instance, what are the legalities of using the data involved? When is a machine learning model truly predictive, or is it merely coincidence? What are the implications of being able to diagnose someone's cancer risk precisely?

Novel sources of information such as wearable devices, trackers, and EHRs place machine learning in the space of innovation and ethical wonderment. The implications for broader healthcare and insurance are massive and should be taken into consideration with machine learning research projects.

As a result, there has been progress in the field of machine learning predictive models, particularly in biomedical research where guidelines for reporting on models have been published, including steps for identifying prediction problems, ensuring n -fold cross-validation, and calibration of outcomes (see Figure 3-11). Calibration has relevance in clinical practice and refers to the relationship between the predictions made by the model and the observed outcomes. As an example, if the prediction model “predicts 70% risk of mortality in the next one year for a patient with lung cancer, then the model is well calibrated if in our dataset, approximately 70% of patients with lung cancer die within the next one year.”

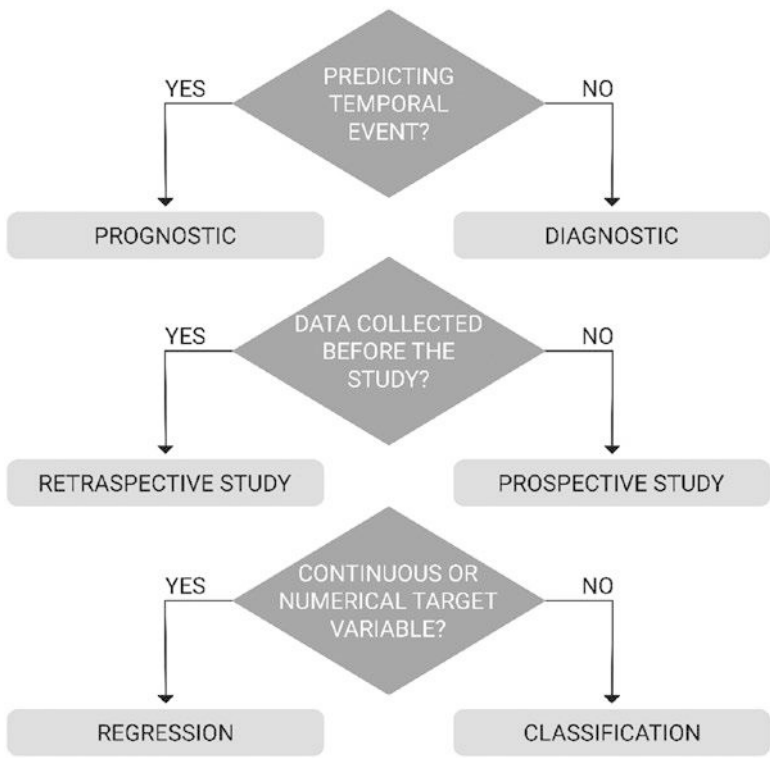


Figure 3-11. Understanding the task at hand Luo W., Phung D., Tran T., Gupta S., Rana S., Karmakar C., Shilton A., Yearwood J., Dimitrova N., Ho T. B., Venkatesh S., and Berk M., “Guidelines for Developing and Reporting Machine Learning Predictive Models in Biomedical Research: A Multidisciplinary View,” *J Med Internet Res* 2016; 18(12):e323, URL: <https://www.jmir.org/2016/12/e323>, doi: 10.2196/jmir.5870, PMID: 27986644, PMCID: 5238707

CHAPTER 4

Machine Learning Algorithms

“Any fool can know. The point is to understand.”

—Albert Einstein

You do not need a background in algebra and statistics to get started in machine learning. However, be under no illusions, mathematics is a huge part of machine learning. Math is key to understanding how the algorithm works and why coding a machine learning project from scratch is a great way to improve your mathematical and statistical skills. Not understanding the underlying principles behind an algorithm can lead to a limited understanding of methods or adopting limited interpretations of algorithms. If nothing else, it is useful to understand the mathematical principles that algorithms are based on and thus understand best which machine learning techniques are most appropriate.

There is an overwhelming number of machine learning algorithms in the public domain. Many are variations (typically faster or less computationally expensive) on several prominent themes. Learners are the nucleus of any machine learning algorithm, which attempts to minimize a cost function—also referred to as an error function or loss function—on training and test data.

Many of your machine learning projects will use defined methods from popular libraries such as numpy, pandas, matplotlib, scipy, scikit-learn, scrapy, NLTK, and so forth.

The most common machine learning tasks are classification, regression, and clustering. Several algorithms can work for functions of discrete prediction, such as k-nearest neighbors, support vector machines, decision trees, Bayesian networks, and linear discriminant analysis.

This chapter provides a comprehensive analysis of machine learning algorithms, including programming libraries of interest and examples of real-world applications of such techniques.

Defining Your Machine Learning Project

Tom Mitchell provides a concise definition of machine learning:

“A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E .” [50]

This definition can be used to aid machine learning projects in assisting us to think clearly about what data is collected and utilized (E), what the task at hand is (T), and how we will evaluate the results (P).

Task (T)

The task refers to what we want the machine learning model to do. The function the model learns is the task. This does not refer to the task of actually learning but rather the task at hand. For example, a robotic Hoover would have the task of hoovering a surface. Often a task is broken down into smaller tasks.

Performance (P)

Performance is a quantitative measure of a machine learning model's ability. Performance is measured using a cost function, which typically calculates the accuracy of the model on the task. The performance can be measured by an error rate, which is the proportion of examples for which the model gives the wrong outputs. The learning method aims to minimize the error rate, ideally without falling into local minima or maxima.

Experience (E)

Experience refers to the amount of labeled data that is available as well as the amount of supervision required by the machine learning model.

Mitchell's reference proves useful in defining machine learning projects. This definition helps understand the data required for collection (E), which develops as the knowledge base; the problem at hand that needs a decision (T); and how to evaluate the output (P).

For example, a prediction algorithm may be assigned task T to predict peak emergency hospital admissions based on experience data E of past hospital admissions and their respective dates and times. The performance measure P could be the accuracy of prediction; and as the model receives more experience data, it would ideally become better in its forecast. It is noteworthy that one criterion does not necessarily determine performance. In addition to accuracy of the prediction, a peak hospital admissions method could provide extra information by predicting peaks that would be the most costly or resource intensive. Table 4-1 gives examples of tasks and the experience which would be created, along with performance metrics.

Table 4-1. *Understanding a learning problem by its components*

Problem	Task	Performance	Experience
Learning how to perform surgical stitching	Stitching a patient's head	Accuracy (perceived pain and/or time could be a feature)	Stitching of patients' heads and feedback on performance measures
Image recognition	Recognizing a person's weight from an image	Accuracy of weight prediction	Training dataset of images of people and their respective weights
Robotic arms used for patient medication delivery	Move the correct medication into the correct patient packet	Percentage of correctly placed medications	Training examples, real-life experience
Predicting the risk of disease	Diagnosing likelihood of type 2 diabetes	Percentage of correctly diagnosed patients	Training dataset of patient health records, real-life experience in prediction and feedback

Machine learning is used to answer questions such as “Is it likely I have type 2 diabetes?”; “What object is this in this image?”; “Can I avoid the traffic?”; and “Will this recommendation be right for me?”

Machine learning is employed favorably in each of these problem domains, yet the complexity of real-world problems means that it is currently infeasible for a specialized algorithm to work faultlessly every time, in every setting, for everyone—whether in healthcare or any other industry.

A common aphorism in math comes from British statistician George Box: “All models are wrong, but some are useful.” The purpose of machine learning is not to make perfect guesses, as there is no such thing as an ideal model. Instead, machine learning builds on foundations in statistics to provide useful, accurate predictions that can be generalized to the real world.

When training a machine learning algorithm, the training dataset must be a statistically significant random sample. If this is not the case, there is a risk of discovering patterns that don’t exist, or the noise. The memorization of training data is known as overfitting. The model remembers the training data, performing well in training but not with previously unknown data. Equally, if the training data is too small, the model can make inaccurate predictions.

The goal of machine learning, whether supervised or unsupervised is generalization—the ability to perform well based on past experiences.

Common Libraries for Machine Learning

Python is increasingly popular within the data science and machine learning industry. All of the libraries are open source on Github, which also provides metrics for popularity and robustness.

- Github is a service based on the Git version control system. The Git approach to storing data is stream-like. Git takes a snapshot of your files and stores references to these snapshots. Git has three states: committed, modified, and staged—which refers to data storage in your local filesystem, local file changes, and flagged files to be committed, respectively. The benefits of using metadata are evident when using Github.

- Numerical Python, or numpy, is an essential Python extension library. It enables fast, N-dimensional array objects for vector arithmetic, mathematical functions, linear algebra, and random number generation. Basic array arithmetic performance is noticeably slower without this library.
- SciPy is a library for scientific processing that builds on numpy. It enables linear algebra routines as well as signal and image processing, ordinary differential equation methods, and special functions.
- Matplotlib is a Python library of similar data plotting and visualization techniques. MATLAB is a standard tool used for data processing. However, the simplicity and usefulness of Matplotlib helps Python to become a viable alternative to MATLAB.
- Pandas is a tool for data aggregation, data manipulation, and data visualization. Within pandas, one-dimensional arrays are referred to as Series, with multidimensional arrays referred to as DataFrames.
- Scikit-learn contains image processing and machine learning techniques. This library is built on SciPy and enables algorithms for clustering, classification, and regression. This includes many of the algorithms discussed in this chapter such as naïve Bayes, decision trees, random forests, k-means, and support vector machines.

- NLTK, or Natural Language Toolkit, is a collection of libraries that are used in natural language processing. The NLTK enables the foundations for expert systems such as tokenization, stemming, tagging, parsing, and classification—which are vital for sentiment analysis and summarization.
- Genism is a library for use on unstructured text.
- Scrapy is an open source data mining and statistics designed initially for scraping web sites.
- TensorFlow is a Google Alphabet-backed open source library of data computations optimized for machine learning. It enables multilayered neural networks and quick training. TensorFlow is used in many of Google's intelligent platforms.
- Keras is a library for building neural networks that builds on TensorFlow.

Supervised Learning Algorithms

In many supervised learning settings, the goal is to finely tune a predictor function, $f(x)$, or the hypothesis. The learning comprises of using mathematical algorithms to represent the input data x within a particular domain. For example, x could take the value of the time of day, and $f(x)$ could be a prediction of waiting time at a particular hospital.

In reality, x typically represents multiple data points. Hence inputs are expressed as a vector. For instance, in the preceding example, $f(x)$ takes in input x as the time of the week. In improving the predictor, we could use not only day (x_1), but time (x_2) of day, weather (x_3), place within the queue (x_4), and so on—on the premise the data was available. Deciding which inputs to use is an integral part of the machine learning design process.

Each record i in the dataset is a vector of features $x(i)$. In the case of supervised learning, there will also be a target label for each instance, $y(i)$. The model is trained with inputs in the form $(x(i), y(i))$. The training dataset can be represented as $\{(x(i), y(i)); i = 1, \dots, N\}$, where x represents the input and y the output.

The model has the simple form:

$f(x) = ax + b$, where a and b are constants, and b refers to the noise in the data.

Machine learning aims to find optimal values of a and b such that the predictions $f(x)$ are as accurate as possible. The method of learning is inductive. The model trawls the training dataset to learn patterns in the data. The model induces a pattern, or hypothesis, from a set of examples. The output of a machine learning prediction model differs from the actual value of the function due to bias, noise, and variance, as discussed in chapter 3.

Optimization of $f(x)$ takes place through training examples. Each training example $\{x_1 \dots x_n\}$ has an input value x_{training} and a corresponding output, Y . Each example is processed, and the difference between the known and correct value y (where $y \in Y$) and the predicted value $f(h_{\text{training}})$ is computed. After processing a suitable number of samples from the training dataset, the error rate of $f(x)$ is calculated, and values a and b are used as factors to improve the correctness of the prediction.

This also addresses inherent randomness or noise in the data, also referred to as irreducible error.

There are two different main approaches when it comes to feature selection. First, an independent assessment based on the general characteristics of the data. Methods belonging to this approach are called filter methods, as the feature set is filtered out before model construction. The second approach is to use a machine learning algorithm to evaluate different subsets of features and finally select the one with the best

performance on classification accuracy. The latter algorithm will be used in the end to build a predictive model. Methods in this category are called wrapper methods because the arising algorithm wraps the entire feature selection process.

The law of large numbers theorem applies to machine learning prediction. The theorem describes how the result of a large number of identical experiments will eventually converge, or even out.

This process is iterated until the system converges on the best values for a and b . As such, the machine learns through experience and is ready for deployment. Typically, we are concerned with minimizing the error of a model or making the most accurate predictions possible.

Within supervised machine learning, there are two main approaches: regression-based systems and classification-based systems.

Classification

Classification is the process of determining a category for an item (or items, x) from a set of discrete predefined categories (or labels, y) through approximating a mapping function, f . Classification algorithms use features (or attributes) to determine how to classify an item into at least a binary class, or two or more classes.

Training data is supplied as an input vector. The input vector consists of labels for item features. For instance, an item may read the following:

{“HbA1c, 7.8%”, “Diabetic”} in the form { x_{training} , output}

It is from the data that classification algorithms learn to classify. Classification problems require examples to be labeled to determine patterns; hence a model’s classification accuracy is best validated through the correctly classified examples of all predictions made.

Algorithms that perform classification include decision trees, Naïve Bayesian, logistic regression, kNNs (k-nearest neighbors), support vector machine, and so forth.

Common use cases of classification algorithms are the following:

- Suggesting patient diagnosis based on health profile—that is, does patient have retinopathy?
- Recommending the most suitable treatment pathway for a patient
- Defining the period in the day in a clinic that there are peak admissions
- Determining patient risk of hospital readmission
- Classify imagery to identify disease—for example, determining whether a patient is overweight or not from a photograph

Regression

Regression involves predicting an output that is a continuous variable. As a regression-based system predicts a value, performance is measured by assessing the number of prediction errors.

The coefficient of determination, mean absolute error, relative absolute error, and root mean square error are statistical methods used for evaluating regression models.

Regression algorithms include linear regression, regression trees, support vector machines, k nearest neighbor, and perceptrons.

It is often possible to convert a problem between classification and regression. For example, patient HbA1c can also be useful if stored as a continuous value. A continuous value of HbA1c = 6.9% can be classified into a discrete category, diabetic, on rules defined on HbA1c as Diabetic (Set 1): $\text{HbA1c} \geq 6.5\%$, Non-diabetic (Set 2): $\text{HbA1c} < 6.5\%$.

Time series refers to input data that is supplied in the order of time. The sequential ordering of the data adds a dimension of information. Time series regression problems are commonplace.

Common use cases of regression algorithms are the following:

- Estimating patient HbA1c from previous blood glucose readings (time series forecasting)
- Predicting the days before hospital readmission (which may be received as a time-series regression problem)
- Approximating patient mortality risk based on clinical data
- Forecasting the risk of developing health complications
- Forecasting when the office printer will be queued with requests from other employees

Decision trees

Decision trees are flowcharts that represent the decision-making process as rules for performing categorization. Decision trees start from a root and contain internal nodes that represent features and branches that represent outcomes. As such, decision trees are a representation of a classification problem. Decision trees can be exploited to make them easier to understand. Each decision tree is a disjunction of implications (i.e., if-then statements), and the implications are Horn clauses that are useful for logic programming. A Horn clause is a disjunction of literals.

On the basis that there are no errors in the data in the form of inconsistencies, we can always construct a decision tree for training datasets with 100% accuracy. However, this may not roll out in the real world and may indicate overfitting, as we will discuss.

For instance, take the simplified dataset Table 4-2, which can be represented as the decision tree as shown in Figure 4-1:

Table 4-2. Risk of type 2 diabetes

Record	Unhealthy	BMI	Ethnicity	Outcome
Person 1	Yes	>25	Indian	High risk
Person 2	Yes	<25	Indian	Low risk
Person 3	Yes	>25	Caucasian	High risk
Person 4	No	<25	Caucasian	Low risk
Person 5	No	>25	Chinese	Medium risk

Classifying an example involves subjecting the data to an organized sequence of tests to determine the label. Trees are built and tested from top to bottom as so

1. Start at the root of the model
2. Wait until all examples are in the same class
3. Test features to determine best split based on a cost function
4. Follow the branch value to the outcome
5. Repeat number 2
6. Leaf node output

The central question in decision tree learning is which nodes should be placed in which positions, including the root node and decision nodes. There are three main decision tree algorithms. The difference in each algorithm is the measure or cost function for which nodes, or features, are selected. The root is the top node. The tree is split into branches; evaluated through a cost function; and a branch that doesn't split is a terminal node, decision, or leaf.

Decision trees are useful in the way that acquired knowledge can be expressed in an easy to read and understandable format (see Figure 4-1). It mimics human decision-making whereby priority—determined by feature importance, relationships, and decisions—is clear. They are simple in the way that outcomes can be expressed as a set of rules.

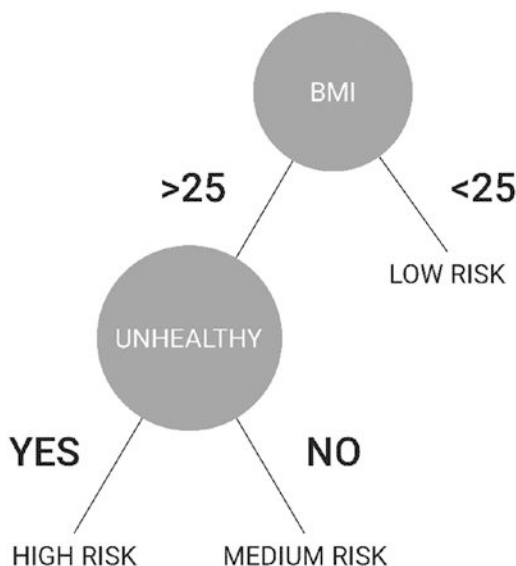


Figure 4-1. *Decision tree of $n = 2$ nodes*

Decision trees provide benefits in how they can represent big datasets and prioritize the most discriminatory features. If a decision tree depth is not set, it will eventually learn the data presented and overfit. It is recommended to set a small depth for decision tree modeling. Alternatively, the decision tree can be pruned, typically starting from the least important feature, or the incorporation of dimensionality reduction techniques.

Overfitting is a common machine learning obstacle, and not limited to decision trees. All algorithms are at risk of overfitting, and a variety of techniques exist to overcome this problem. Random forest or jungle decision trees can be extremely useful in this.

Pruning reduces the size of a decision tree by removing features that provide the least information. As a result, the final classification rules are less complicated and improve predictive accuracy.

The accuracy of a model is calculated as the percentage of examples in the test dataset that is classified correctly.

- True Positive / TP: Where the actual class is yes, and the value of the predicted class is also yes
- False Positive / FP: Actual class is no, and predicted class is yes
- True Negative / TN: The value of the actual class is no, and the value of the predicted class is no
- False Negative / FN: When the actual class value is yes, but predicted class is no
- Accuracy: $(\text{correctly predicted observation}) / (\text{total observation}) = (TP + TN) / (TP + TN + FP + FN)$
- Precision: $(\text{correctly predicted Positive}) / (\text{total predicted Positive}) = TP / TP + FP$
- Recall: $(\text{correctly predicted Positive}) / (\text{total correct Positive observation}) = TP / TP + FN$

Classification is a common method used in machine learning; and ID3 (Iterative Dichotomizer 3), C4.5 (Classification 4.5), and CART (Classification And Regression Trees) are common decision tree methods where the resulting tree can be used to classify future samples.

Iterative Dichotomizer 3 (ID3)

Python: scikit-learn; Method: decision-tree-id3

Ross Quinlan has contributed significantly to the area of decision tree learning, developing both ID3 and C4.5.[51] The ID3 algorithm uses the cost function of information gain, based on work from “the father of information theory” Claude Shannon. Information gain is calculated through a measure known as entropy.

Entropy determines how impure, or disordered, a data set is. Given an arbitrary categorization C , into categories $\{c_1, \dots, c_n\}$ and a set of examples S , for which the proportion of c_i is p_i , the entropy of sample S is the following:

$$H(S) = \sum_{x \in X} -p(x) \log_2 p(x)$$

Information gain is calculated using entropy. This value is given for attribute A , with respect to examples S . The values attribute A can take, $\{t_1, \dots, t_n\}$, representing the totality of set T .

$$IG(A, S) = H(S) - \sum_{t \in T} p(t) H(t)$$

ID3 selects the node with the highest information gain to produce subsets of data that are then applied the ID3 algorithm on recursively. The information gain of an attribute can be understood as the expected loss in entropy as the result of learning the value of attribute A . The algorithm terminates once it exhausts the attributes or the decision tree classifies the examples completely. ID3 is a classification algorithm that is unable to cater for missing values.

For example, if the following row were to be received, the model would categorize the outcomes as high risk. There is a chance that ethnicity is a determining factor in the outcome. However, there is not enough data for the model to learn the potential outcome.

Record	Unhealthy	BMI	Ethnicity	Outcome
Person 5	No	<25	Chinese	?

C4.5

Python: scikit-learn; Method: C45algorithm

C4.5 is an enhancement built on ID3 that takes number and size of branches into account when choosing an attribute. C4.5 can handle continuous and discrete labels and uses gain ratio as its splitting criteria. It improves on the execution time and accuracy as the size of the dataset increases. C4.5 introduces pruning—a bottom-up technique also known as subtree raising—and subtree replacement to ensure the model doesn’t overfit the data.

CART

Python: scikit-learn; Method: cart

CART, or Classification And Regression Trees, use Gini index as the cost function. Represented as a binary tree, as with C4.5, CART can be used for regression and classifications problems. While entropy is used for exploratory analysis, Gini is to minimize misclassification.

Ensembles

An ensemble is a form of machine learning algorithm and is a popular technique employed in the field of machine learning. Ensemble methods use a collection of machine learning models or learners. As a result, a group of weak learning models is combined to create a more accurate learner.

Bagging

Bagging is an ensemble technique that involves creating multiple independent models with datasets using the bootstrap sampling technique (see Figure 4-2). The technique aids in reducing error by reducing variance. In general, bagging improves predictions as anomalous behaviors are averaged out. In comparison, a collection of iterative predictive models is known as boosting.

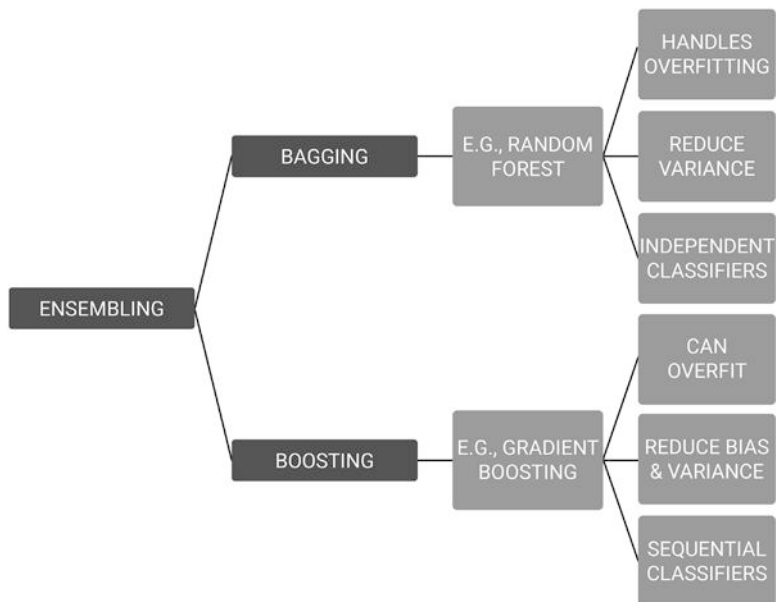


Figure 4-2. *Bagging vs. boosting*

In bagging, $n' < n$ samples are taken from dataset D with replacement. A random selection of features for constructing the best split is chosen each time.

An average, majority vote, or statistical method, is used to determine the average prediction from a model. Random forest decision trees (see Figure 4-3) extend the concept of bagging, as they use a random subset of features too, resulting in less correlation among predictions from learners.

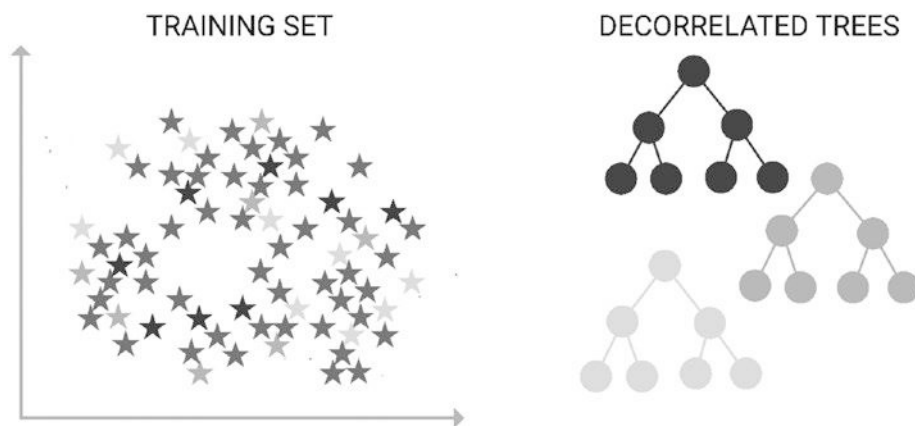


Figure 4-3. *Random forest decision trees*

Random Forest Decision Trees

Python: scikit-learn; Method: RandomForestClassifier

Random forests are created through training multiple decision trees at the same time and refer to a collection of decision tree learners (hence the “forest”). The purpose of random forests is to prevent overfitting. The more decision trees in the random forest, the more accurate results are. Each random forest takes a sample of the dataset and a random subset of features for decision.

Random forest decision trees can be used for classification and regression problems. With regression problems, an average of the outcomes is taken; whereas with classification problems, it is the majority

category. A random forest decision tree is an ensemble of different models. This is an ensemble method due to the procedure reducing variance in the algorithm through combining predictions from multiple models.

The random forest model seems to work particularly well when data is limited to classification or decision problems by bootstrapping existing data. A benefit of random forest classifiers is that it can also handle missing values.

Boosting

Boosting is an ensemble method that produces a collection of predictive models iteratively. The concept involves each new model learning from the errors of previous models. Boosted methods are typically implemented when there is a plethora of data for prediction.

Through iteratively creating new models to compensate for error, the process “boosts” or converts weak learners into strong learners—reducing their bias and variance. Each iteration of the technique learns relationships in the data and is then subsequently analyzed for errors, with incorrect classifications given a weighting. The aim is to minimize the error from the previous model. The underlying machine learning method used for boosting algorithms can be anything.

Boosting works through the following process:

1. Data is represented
2. Decision stump is made on the most significant cut of features
3. Weightings are given to misclassified observations
4. The process is repeated and all stumps combined to achieve final classifier

Gradient Boosting

Python: scikit-learn; Method: GradientBoostingClassifier, GradientBoostingRegressor

Gradient boosting is known by many names including multiple additive regression trees, stochastic gradient boosting, and gradient boosting machines. Gradient boosting is a technique derived from decision trees when Jerome Friedman applied the gradient boosting concept to decision trees in 2001.[52] AdaBoost (Freund and Shapire, 2001) is a gradient boosting variant. In gradient boosting, the model is trained sequentially. Each iterative model seeks to minimize the loss function.

Similar to random forest decision trees, gradient boosting is made from weak predictors. Each new tree is subsequently trained with the data previously incorrectly classified by the previous tree. This iterative procedure expedites the model by focusing on more challenging data as easier predictions are made early on in learning.

XGBoost, or extreme gradient boosting, are variations of gradient boosting. XGBoost adds regularization and takes advantage of computing power with distributed, multithreaded processing for greater speed and efficiency.

Adaptive Boosting

Python: scikit-learn, numpy; Method: AdaBoostClassifier, AdaBoostRegressor

Adaptive boost is a popular method that iteratively learns from mistakes. Decision tree stumps are weak learners, splitting data by one rule only. The decision tree stump is improved by focusing on samples that were incorrectly classified. These are identified through weightings that are associated with the data. The model can learn from its mistakes and provides a final solution with lower bias than a single decision tree.

Linear Regression

Python: scikit-learn; Method: `linear_model.LinearRegression`

Linear regression is a continuous statistical technique to understand the relationship between input (x: predictors, explanatory or independent variables) and output (y: response or dependent variable). As the name suggests, linear regression (see Figure 4-4) requires that y be calculated from a linear combination of input variables. The aim is to place a line of best fit through all of the data points, which enables easy understanding of predictions by minimizing margin or residual variation.

$$\min \sum (y_i - \hat{y}_i)^2,$$

where y_i is the actual observed response value; \hat{y}_i is the predicted value of response variable; and their subtraction is the residual variation or difference between observed and predicted values of y.

- Criterion variable: outcome variable we are predicting
- Predictor: what variable predictions are based on
- Intercept: where the line of best fit intercepts the y-axis

When there is one input variable x, this technique is known as simple linear regression. Multiple linear regression evaluates two or more predictor variables.

$$Y = \beta_0 + \beta_1 X$$

$$\beta_1 = \frac{\sum_{i=1}^m (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^m (x_i - \bar{x})^2},$$

$$\beta_0 = \bar{y} - \beta_1 \bar{x}$$

where y is the output variable, x is the predictor, B_0 and B_1 are coefficients that are estimated to move the line of best fit, \bar{x} refers to the mean of x , B_0 is the bias or intercept (where it intercepts the y -axis), and B_1 is the slope or rate at which y increases/decreases per unit of x .

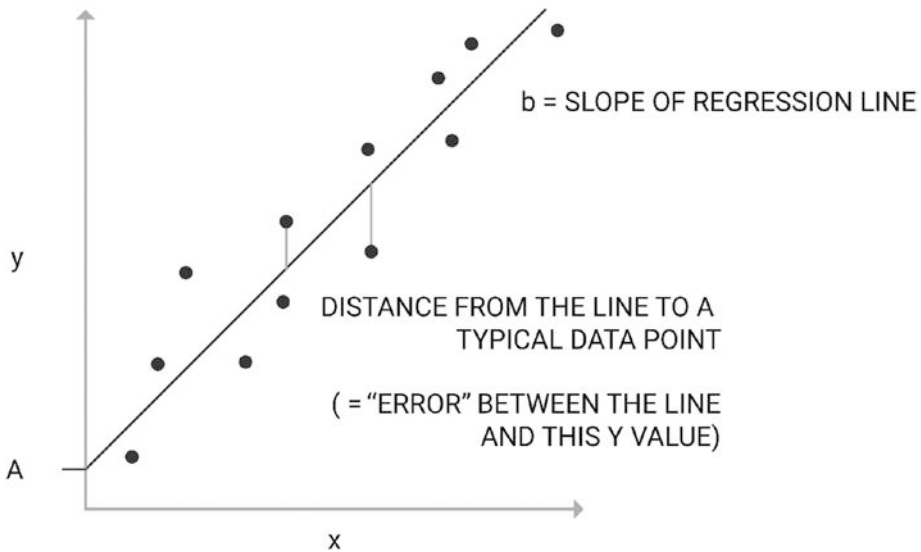


Figure 4-4. Linear regression

Once the model has learned the parameters, they are used to predict values of y given new, unseen input X .

Popular techniques used to estimate the coefficients are ordinary least squares and gradient descent. Given a line of best fit placed through a sample dataset with a linear relationship, ordinary least squares attempts to minimize the sum of the squared distances from each data point to the regression line. The intention is to minimize the sum of squared residuals. Coefficient values are iteratively optimized to reduce the error of the model.

Ordinary least squares start with random values for B_0 and B_1 . The sum of squared errors is calculated for each (input, output) pair. A learning rate is chosen, which acts as a scale factor; and the coefficients are optimized toward minimizing the error until no further gains can be made.

A learning rate is a hyperparameter that has its value determined by experimentation. Different values are tried, and the value that gives the best results is used.

Criticisms of linear regression focus on its simplicity, resulting in it being unable to capture complex relationships within the data. Noise in the data can also be learned by the model, resulting in overfitting. Relationships between variables are not always linear, and therefore a linear regression model can perform poorly on data with nonlinear relationships. However, variables can sometimes be transformed to fit a linear regression model. Linear regression also assumes variables are homoscedastic, which means all variables in the vector have the same variance.

The aforementioned is a perfect demonstration as to why when training a model, one should use a variety of machine learning algorithms.

Logistic Regression

Python: scikit-learn; Method: `linear_model.LogisticRegression`

Linear regression applies to continuous variables, whereas logistic regression predictions apply to discrete values or classification problems after the application of a transformation function.

Developed by statistician David Cox in 1958, logistic regression (or logit regression) is often used for binary classification problems. For example, logistic regression can be used to predict whether a patient will experience an adverse event given a particular medication; whether a patient is likely to be readmitted to a hospital or not; or whether a patient has a particular disease diagnosis.

Unlike linear regression, the output of the model comes in the form of a probability between 0 and 1. The predicted output is generated by log-transforming the x input and using the logistic function $h(x) = 1/(1 + e^{-x})$. A specified threshold is used to transform this probability into a binary classification.

Logistic regression uses the sigmoid function, an S-shaped curve (see Figure 4-5) that can take any continuous value and map it into a probability value between 0 and 1 for a default class.

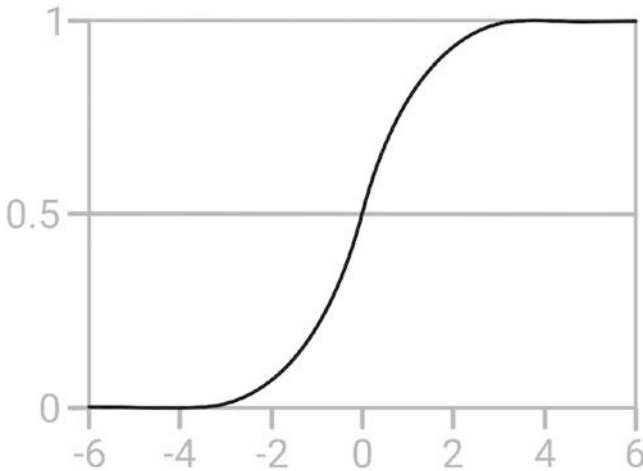


Figure 4-5. *Logistic function*

Sigmoid function: $1/(1 + e^{-\text{value}})$, where e Euler's number and value is the value requiring transformation.

Logistic regression uses an equation $P(x) = e^{(b_0 + b_1 \cdot x)} / (1 + e^{(b_0 + b_1 \cdot x)})$, which can be transformed into $\ln(p(x)/1-p(x)) = b_0 + b_1 \cdot x$. As with linear regression, b_0 and b_1 are values learned from the training data.

As an example, take a model that determines whether a growth on the skin surface is benign or not. The input variable, x , could be the size, depth, or texture of the growth; and the default variable $y = 1$ indicates a benign growth. As depicted, the logistic function transforms the x -value of the various instances into a probability value between 0 and 1. If the probability crosses the threshold of 0.5, the tumor is classified as benign.

Logistic regression seeks to learn values of b_0 and b_1 through training that minimize the error between predicted and actual outcomes. Maximum likelihood estimation is an iterative procedure used for this purpose. Maximum likelihood estimation starts with a random value as the

best weight for each predictor variable and then adjusts these coefficients iteratively until there is no improvement in the ability to predict the outcome value. The final coefficients have a value that minimizes the error in the probabilities predicted.

As with linear regression, logistic regression assumes a linear relationship between input variables and output. Feature reduction may be required to transform data into linear models. Logistic regression models are also prone to overfitting, which may be overcome through the removal of highly correlated inputs. Finally, it is possible that coefficients fail to converge, which can happen if data is sparse or highly correlated.

SVM

Python: scikit-learn Method: `svm.SVC`, `svm.LinearSVC`

Support vector machine (SVM) is a nonprobabilistic binary linear classifier used with both classification and regression problems. SVM is often used along with NLP methods to analyze text for topic modeling and sentiment analysis. It is also used in image recognition problems and handwriting digit recognition.

The algorithm finds a hyperplane, or line of best fit, between two classes determined by the support vectors. Support vectors are the data points nearest the hyperplane that would alter the position of the hyperplane if removed. The greater the margin value, or distance from the data points to the hyperplane, the more confidence there is that the data is appropriately classified. The line of best fit is learned from an optimization procedure that seeks to maximize the margin.

SVM uses something known as kernelling to map data to higher dimensional feature spaces. Data is mapped using the kernel trick in iteratively more dimensions until a hyperplane can be formed to classify it.

Take the problem in Figure 4-6. On the left, we see that it would be impossible to find a line of best fit in two dimensions. SVM uses the kernel trick to map data into three dimensions and can define a hyperplane to classify the data.

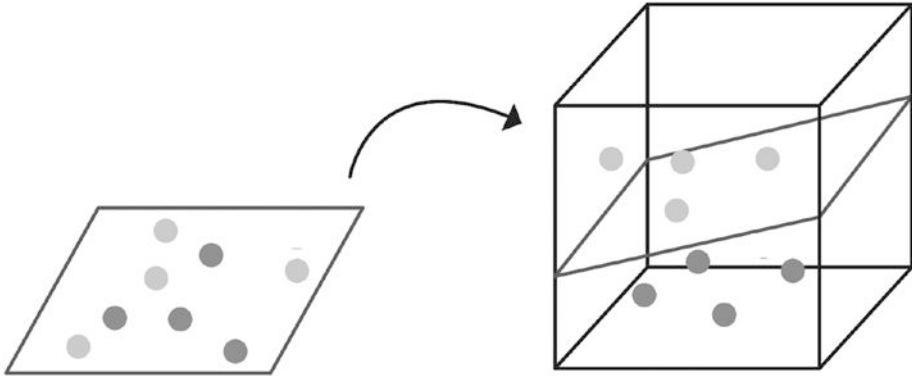


Figure 4-6. SVM visualization

SVMs take numerical inputs and work well on small datasets. However, as dimensionality increases, the ability to understand and explain the model reduces. As dataset size increases, so too can training time. SVMs are also less capable on noisy data.

The SVM can be expressed as a sum over the support vectors (see Figure 4-7):

$$f(x) = \sum_i \alpha_i y_i (\mathbf{x}_i^T \mathbf{x}) + b$$



 SUPPORT VECTORS

Figure 4-7. SVM expressed as a sum over the support vectors

Naive Bayes

Python: scikit-learn; Method: GaussianNB, MultinomialNB, BernoulliNB

Naïve Bayes uses Bayes' theorem to calculate the probability an event will occur given another event has already occurred. The algorithm is acknowledged as naïve because it assumes all variables are independent of each other, which is atypical of real-world examples. The Bayesian classifier is often used when input dimensionality is high.

For example, given a specified threshold, this method can be used to probabilistically classify whether a vector is in one class or another:

$$P(B|A) = (P(A|B) * P(B))/P(A)$$

- $P(B|A)$ = Posterior probability (see Figure 4-8).
Probability of hypothesis B being true, given the data A,
where $P(B|A) = P(a_1|b) * P(a_2|b) * \dots * P(a_n|b) * P(A)$
- $P(A|B)$ = Likelihood: the probability of data A given that the hypothesis B was true
- $P(A)$ = Class prior probability
- $P(B)$ = Predictor prior probability

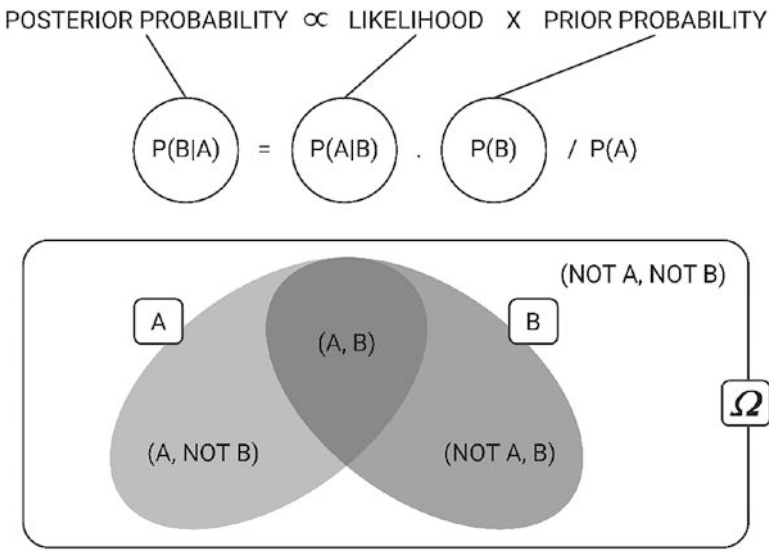


Figure 4-8. Posterior probability

The prior distribution $P(B)$ and likelihood probability $P(B|A)$ can be estimated from the training data. To calculate which class an example belongs to, the probability of being in each class is calculated. The class assigned to an example is the class that produces the highest probability for the example.

Table 4-3. Outcome liklihood

Record	Unhealthy	Outcome
Person 1	Yes	High risk
Person 2	Yes	Low risk
Person 3	Yes	High risk
Person 4	No	Low risk

Using Table 4-3 as a simple example, we can learn the probability of a patient's disease risk from whether they are unhealthy or not. To determine the outcome for unhealthy = yes, we calculate $P(\text{high risk}|\text{unhealthy})$ and $P(\text{low risk}|\text{unhealthy})$ and choose the outcome with the highest probability.

$$\begin{aligned} &P(\text{high risk}|\text{unhealthy}) \\ &= (P(\text{high risk}|\text{unhealthy}) * P(\text{high risk})) / P(\text{unhealthy}) = \\ &\quad (2/2 * 2/4) / (3/4) = 0.66 \end{aligned}$$

Naïve Bayesian models are easy to build and particularly useful for extensive data sets. Along with its simplicity (and hence its naivety), Naïve Bayes is known to outperform even highly sophisticated classification methods.

kNN: k-nearest neighbor

Python: scikit-learn; Method: neighbors.KNeighborsClassifier

Not to be confused with k-means clustering, the kNN method classifies an unknown object O with the label of the majority of the k -nearest neighbors. Used in classification and regression problems, kNNs are a nonparametric technique; they do not learn a model. Instead, kNNs store the training dataset as its representation and perform classification of a new sample based on learning by analogy. As there is no learning of the model, kNNs are considered lazy learners.

Each object represents a point in N -dimensional space. A neighbor is defined nearest if it has the smallest distance in feature space. The distance between an unseen object and its neighbor is calculated using the Euclidian distance between $A = (a_1 \dots a_n)$ and $B = (b_1 \dots b_n)$ as follows:

$$d(A, B) = \sqrt{\sum_{i=1}^n (a_i - b_i)^2}$$

The algorithm works as such:

1. Calculate the distance between any two points
2. Find the nearest neighbors based on these pairwise distances
3. Majority vote on a class label based on the nearest neighbors list

The prediction is made at the time of the request. In regression problems, the mean or median of the k -most similar instances is used. In classification problems, the class with the highest frequency from the k -most similar instances is selected. As with most algorithms, the determinant method can vary. The distance between binary vectors (Hamming distance) and the sum of absolute difference between real vectors (Manhattan distance) are also used.

A disadvantage of kNNs is the large computing requirement for classifying an object, as the distance for all neighbors in the training dataset must be calculated. It is not particularly well suited to high-dimensional data. Each predictor variable can be considered a dimension of n -dimension input space. For instance, x_1 would be one-dimensional and x_1, x_2 two-dimensional, and so on. The increase in dimensionality exponentially increases the volume of the input space.

kNNs are not well suited to data with missing values, as distances between vectors cannot be calculated on missing data.

Neural Networks

The biological mechanism in which the brain works inspires artificial neural networks, a different paradigm for computing, based on the parallel architecture of animal brains.

Neural networks are well-equipped for data mining tasks due to their ability to model multidimensional data and efficiency at finding hidden

patterns among data. Neural networks can be applied to prediction and classification problems. The process of estimating the result and comparing it with the real output is known as forward propagation.

Neural networks are a subset of machine learning algorithms, which includes perceptrons, fully connected neural networks, convolutional neural networks, recurrent neural networks, long short-term memory neural networks, autoencoders, deep belief networks, generative adversarial networks, and many more. Most of them are trained with an algorithm called backpropagation.

Artificial neural networks (ANNs) are used in a plethora of tasks:

- Classifying disease: computers can learn what images of diseased organs—such as kidneys, eyes, liver—look like to predict the likelihood of disease
- Recognizing speech
- Translating and digitalizing text
- Facial recognition

Neural networks can learn how to accomplish a task like a human brain—supervised, unsupervised, and through reinforcement learning.

Perceptron

Perceptrons are the fundamental units of ANNs. In AI, the perceptron is a node or unit that takes multiple inputs and has a single output, typically 1 or -1. Perceptrons take a weighted sum of the inputs, S , and a unit function calculates the output for the node (see Figure 4-9). Unit functions can be the following:

- Linear—such as the weighted sum
- Threshold functions that only fire if the weighted sum is over a threshold

- Step functions that output the inverse, typically -1, if S is less than the threshold
- Sigma function ($1/(1 + e^{-s})$), which allows for backpropagation

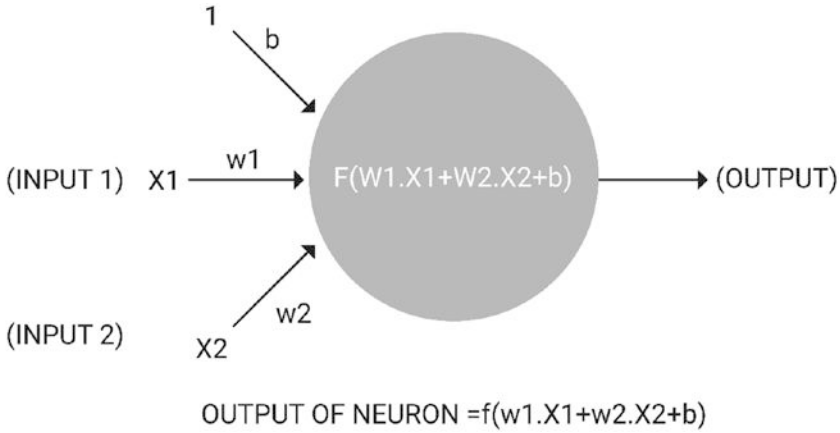


Figure 4-9. Configuration of an artificial neuron

Perceptrons mimic biological neurons (see Figure 4-10), where dendrites receive an input. If the signal is high enough within a set duration of time, the neuron sends an electrical pulse to the terminals then received by dendrites of other neurons. Each perceptron has a bias, similar to b of the linear function $y = ax + b$. It moves the line up and down to fit the prediction with the data better.

$$a = f\left(\sum_{i=0}^N w_i x_i\right),$$

where 1 is x_0 and $b = w_0$.

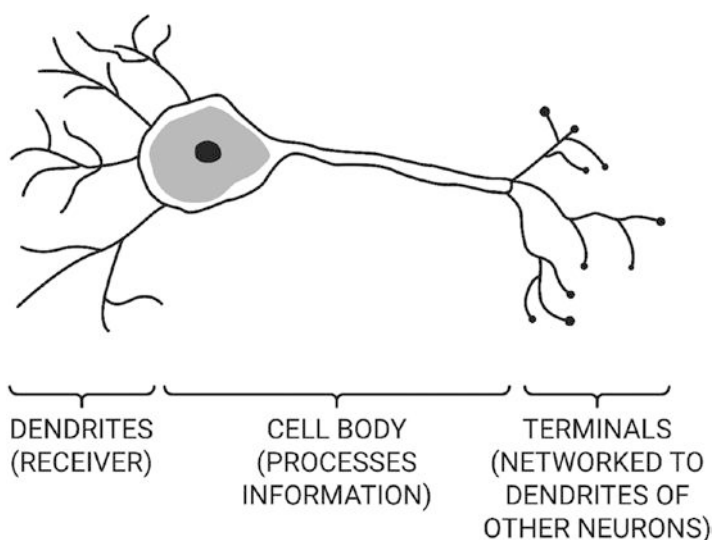


Figure 4-10. *Configuration of a biologic neuron*

Artificial Neural Networks

Python: scikit-learn; Method: linear_model.Perceptron

Artificial neural networks are made of perceptrons and contain one or more hidden layers (see Figure 4-11). There are several topologies, with the simplest being a feedforward network. Backpropagation is the method of determining the error, or loss, at the output and propagating it back to the network. Weights at each node are updated to minimize the respective error output from each neuron. Backpropagation seeks to minimize the error for neurons and thereby the entire model through updating respective neuron weights.

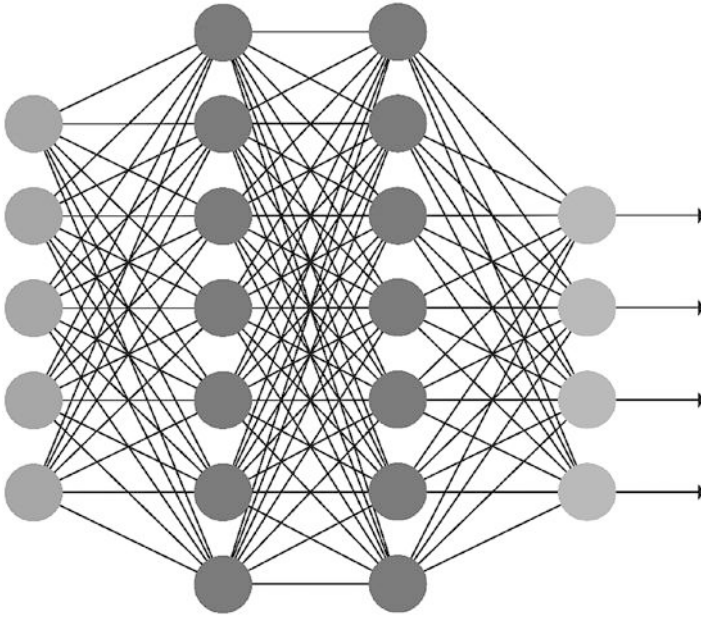


Figure 4-11. *Organization of an artificial neural network*

The primary learning that must be accomplished in neural network models is training neurons when to fire. An epoch refers to one training iteration of forward- and backpropagation. During the learning phase, nodes in a neural network adjust their weights based on the error of the last test result. The learning rate controls how fast the network learns.

The number of perceptrons or neurons in a model are equal to the number of variables in the data. Some representations may also have a bias node. Typically, artificial neural networks have neurons organized in layers. Each layer can perform different transformations on the input data it receives.

Deep Learning

Deep learning is a process that employs deep-layered neural network architectures. Training a deep, artificial neural network model requires more time and CPU power compared to other types of models. Noteworthy

is that the performance of ANN models may not necessarily be superior to typical supervised learning techniques.

Deep learning is not a new technique; however, it is becoming more frequent with hardware advances, primarily power and cost, which have made such computing feasible. Deep neural networks (DNN) were used by Alphabet to create the strongest Go player in history, defeating several human champions of Go in 2017.[53] The ancient game of Go has historically been considered a challenge for AI and demonstrates the capabilities of deep learning (see Figure 4-12).

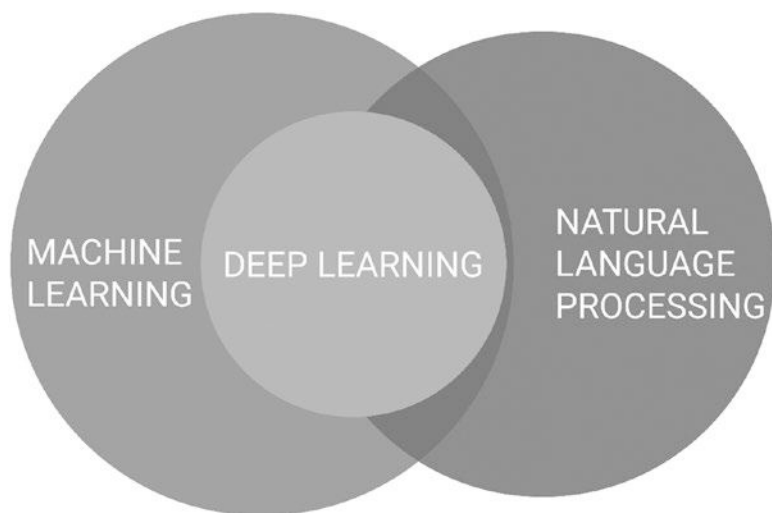


Figure 4-12. *Deep learning: Where does it sit?*

There are many types of artificial neural networks, including the following.

Feedforward Neural Network

In a feedforward neural network, the sum of the weighted inputs is fed to the output (see Figure 4-13). The output is activated (typically 1) when the sum exceeds the unit function threshold. If it doesn't fire, it typically outputs -1.

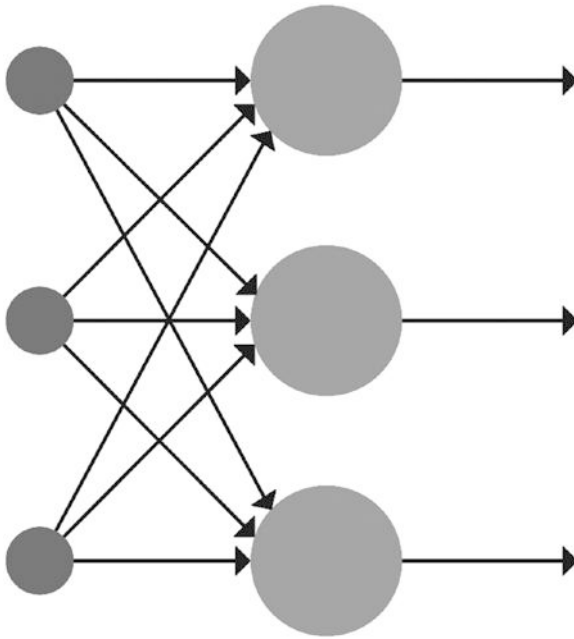


Figure 4-13. *Feedforward neural networks*

Recurrent Neural Network (RNN)—Long Short-Term Memory

Recurrent Neural Networks or RNNs save the output of a layer and feed it back to the input to improve predictions of the output layer (see Figure 4-14). As such, each node has a memory when performing computations and makes use of sequential information. As a result, nodes are deemed to have a memory.

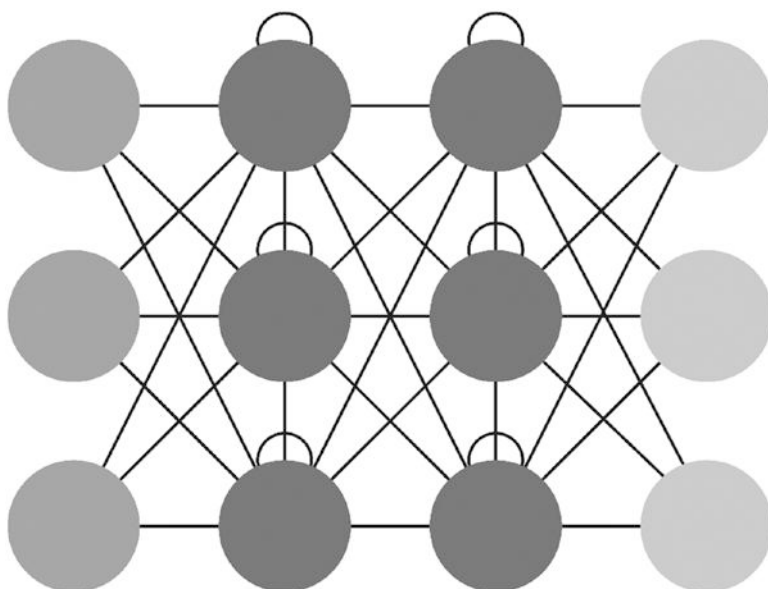


Figure 4-14. *Recurrent neural networks*

Convolutional Neural Network

Convolutional neural networks are deep, feedforward neural networks. Layers are known as convolutional layers and are often used in speech recognition, spatial data, NLP, and computer vision problems. Separate convolutional layers typically apply to different aspects of a particular problem. For instance, in understanding whether there is a face present within a photograph, separate convolutional layers within a DNN may be used to identify different aspects of the face—the eyes, nose, ears, mouth, and so forth.

Modular Neural Network

Similar to the human brain, the concept of modular neural networks is that of neural networks working together.

Radial Basis Neural Network

A neural network that uses radial basis activation functions where the value depends on the distance from the origin (see Figure 4-15).

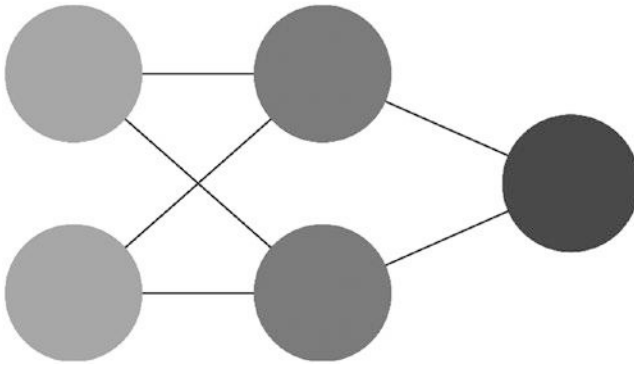


Figure 4-15. *Radial basis neural network*

The key strength of ANNs and deep learning lie in the ability to generate complex linear and nonlinear models from the training dataset.

DNNs (deep and multilayered ANNs) can learn complex relationships between data from training data alone, which typically improves generalization. DNNs can still succumb to the problem of overfitting data, so appropriate pruning may be required. As the name suggests, deep networks can also take a long time to train and require a large number of training examples.

A common concern with neural networks, particularly in healthcare, is that they act as a black box and do not present a clear interface for a user to understand findings. Practically, an ANN model is difficult to interpret, and weights and biases aren't readily interpretable to what is important in the model.

It is typically considered that for supervised learning problems, the upper bound on neurons to prevent overfitting is

$$N_h = N_s(\alpha * (N_i + N_o)),$$

where

N_i = number of input neurons;

N_o = number of output neurons;

N_s = number of samples in training dataset; and

α = an arbitrary scaling factor.

Unsupervised Learning

Unsupervised learning refers to the process of learning a model from unlabeled data. This means that input data (x) is supplied without output (y). Algorithms are left to determine notes within data, which means there is no right or wrong answer.

Semi-supervised learning occurs when some output labels (y) are supplied. For example, learning would be semi-supervised in a model learning to predict diabetic retinopathy from patient eye scans with partially labeled data. If only a subset of eye scans had appropriately labeled outputs (e.g., retinopathy, not retinopathy), the model might not have enough appropriate data to learn a model. As more labeled data is provided for the model to train on, the system would likely become more accurate.

Unsupervised learning can be resource consuming concerning time, money, and expertise. Usually, vast improvements in model accuracy can be made through improving labeling of data.

Unsupervised learning is composed of two main problem concepts: clustering and association.

Clustering

Clustering refers to the process of discovering relationships within the data. Clustering is used for a variety of healthcare uses including the following:

- Grouping patients of similar profiles together for monitoring
- Detecting anomalies or outliers in claims or transactions
- Defining treatment groups based on medication or condition
- Detecting activity through motion sensors

K-Means

Python: scikit-learn; Method: cluster.KMeans

K-means clustering (see Figure 4-16) aims to find k similar groups within the data. K-means is an iterative algorithm that calculates the centroids of the defined k clusters, with all training data assigned to a cluster. Data points are assigned by the distance between the data point and its centroid. A cluster centroid is a collection of attribute values that define the resulting clusters.

$$\operatorname{argmin}_{c_i \in C} \operatorname{dist}(c_i, x)^2,$$

where c_i is the collection of centroids in set C, and distance (dist) is standard Euclidean distance.

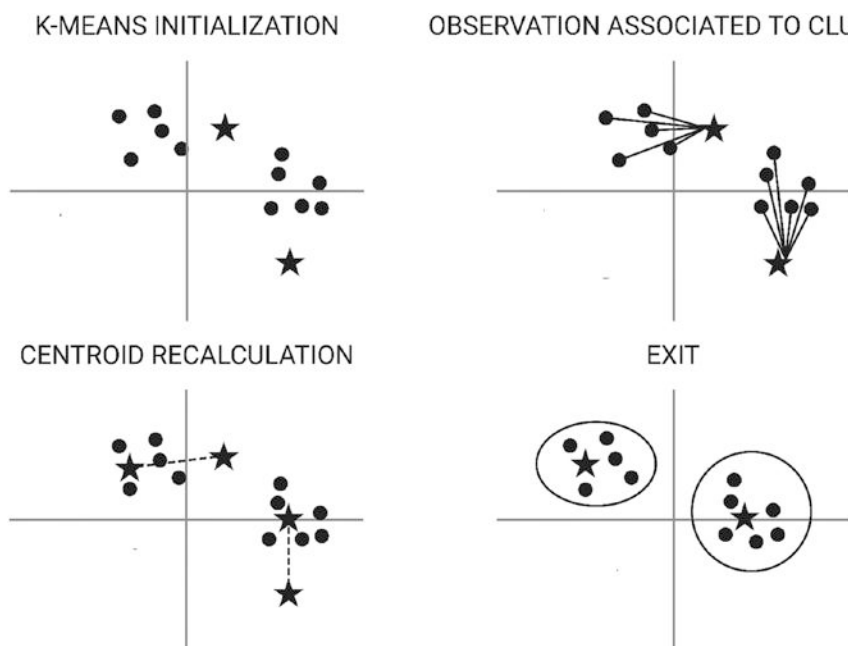


Figure 4-16. *K-means clustering*

Rather than defining groups before data exploration, clustering enables the model to determine groups that have formed organically.

Steps of the K-means algorithm include the following:

- Step 1: Set a value of k . Here, let us take $k = 2$.
 - Assign each data point randomly to any of the $k = 2$ clusters.
 - Determine the centroid for each of the clusters.
- Step 2: For each data point, associate it to the closest cluster centroid.
- Step 3: Recalculate centroids for the new clusters.

Step 4: Continue

- this process until the centroid clusters remain unchanged.

Association

Association rule learning methods extract rules that best explain perceived relationships between variables in data. These rules can discover useful associations in large multidimensional datasets that can be used to drive and optimize.

Association rule learning is historically best applied to online shopping checkout basket datasets gathered on users' purchasing habits. For instance, when someone buys a pizza, they may also tend to purchase wine, just like how someone who purchases lettuce may tend to buy tomatoes, cucumbers, and onion too. Through analyzing transactional datasets, the probability of associations can be predicted. We know that certain items (whether that be food, clothes, or even disease) frequently occur together, and association rule learning seeks to understand these relationships.

In healthcare, in particular, associative symptoms can be understood to predict better and diagnose disease and adverse events. Potential adverse effects based on medication and associative patient comorbidity pathways could lead to improved care and treatment pathways.

There are three important metrics to be familiar with:

Support

Support is the value of absolute frequency. An association rule holds with support sup in dataset T if the $\text{sup} \%$ of transactions contain $X \cup Y$. This represents how popular an itemset is, as measured by the proportion of transactions in which an itemset appears.

$$\text{sup} = \Pr(X \cup Y) = \text{count}(X \cup Y) / \text{total transaction count}$$

Confidence

The confidence measure represents correlative frequency. An association rule holds in dataset T with confidence conf if the $\text{conf} \%$ of transactions that contain X also contain Y. This estimates how likely item Y is to occur or be present in the transactional dataset when item X occurs. This is expressed as $\{X \rightarrow Y\}$ and measures the proportion of transactions with item X in which item Y also appears.

$$\text{conf} = \Pr(Y|X) = \text{count}(X \cup Y) / \text{count}(X)$$

Lift

Lift determines how likely item Y is given that X occurs while accommodating for Y's popularity.

$$\text{Lift} = \text{Support}(X \cup Y) / \text{Support}(X) * \text{Support}(Y)$$

Association rules determine the probability of Y given that X occurs. The benefits of these associations have applications across industry. There are many association rule mining algorithms that use different strategies toward understanding associations and use different structuring of data.

Apriori

Apriori is the most popular association rule mining algorithm. The Apriori algorithm is frequently used in pattern mining and in the identification of features and items that occur together. It is used extensively in basket analysis and is a powerful tool to find hidden feature patterns. It is usually applied to transactional databases to mine frequent relationships and generate association rules.

Apriori begins with a set of n items. The algorithm calculates all possible candidate itemsets, or frequent itemsets, which meet the specified threshold support and confidence values. The Apriori property states for an itemset to be frequent, all of its subsets must also be frequent.

The Apriori algorithm uses a bottom-up approach; frequent subsets are identified and increased one item at a time.

The Apriori method is relatively simple (see Figure 4-17):

- Apply minimum support to find all frequent sets of k items within dataset T where a frequent itemset has support $\geq \text{min}(\text{sup})$
- Expand selection to find frequent sets of $k + 1$ items using frequent k -itemset

The maximum size of an itemset is k , the number of items.

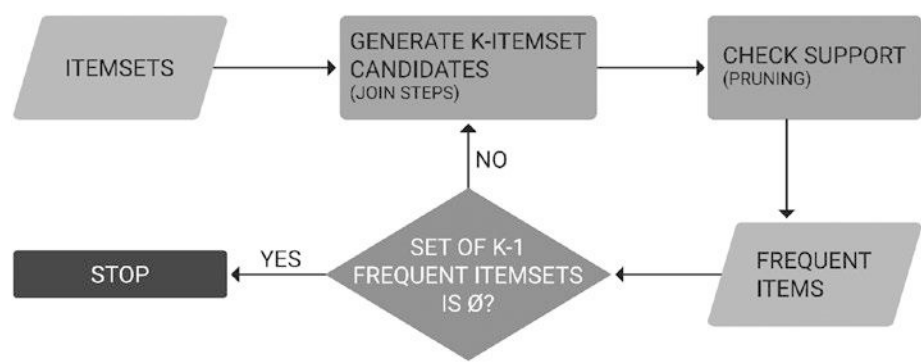


Figure 4-17. *Apriori*

Dimensionality Reduction Algorithms

Although the message is broadly the more data, the better, datasets can often contain many variables and as a result make capturing the signal of the data a more laborious task. Whether it's sparse values, missing values, identifying relevant features, resource efficiency, or more straightforward interpretation, dimensionality reduction algorithms are very useful to data scientists.

As the name suggests, dimensionality reduction algorithms or DRAs reduce the number of dimensions that exist within a dataset. The need for dimensionality reduction becomes apparent when you consider the wealth of data available to make narrow decisions:

- Mobile phones collect hundreds of data points including calls, texts, steps, calories burned, floors climbed, Internet usage, and so forth. What data is best for understanding phone usage?
- Brands on social media are collecting data on engagement and interactions such as comments, likes, followers, sentiment, and mood. What data is best for understanding attitudes toward health?
- Medical health records contain a wealth of information, but only some information is relevant in predicting disease risk or illness progression. What data is relevant to understanding future risk of disease or adverse event?

Dimensionality reduction refers to converting a dataset of many dimensions into fewer dimensions while concisely representing similar data.

The graph in Figure 4-18 shows the relationship between cm and inches for the dataset. This two-dimensional representation can be transformed succinctly into a one-dimensional representation, improving the visualization for understanding. N dimensions of data can be reduced to k dimensions where $k < N$. Dimensions of value can be identified, combined, or new dimensions created to represent the inherent relationships.

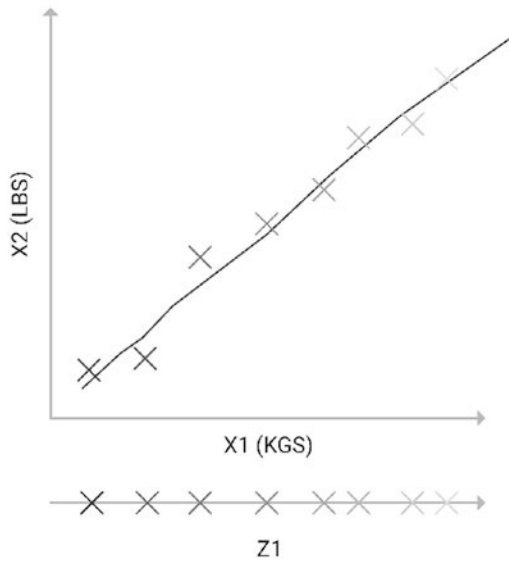


Figure 4-18. *Dimension reduction*

Dimension reduction is extremely useful in machine learning tasks, with a plethora of benefits:

- Fewer dimensions result in quicker computations when compared to the original dataset.
- By default, dimension reduction algorithms reduce the space required for storage.
- Reducing data into less than three dimensions enables visualization and easier understanding.
- Redundant data is removed, which improves the performance of the machine learning model.
- Noise is removed, which improves model performance.

Dimension Reduction Techniques

Dimension reduction can be achieved in a variety of ways:

Missing/Null Values

Missing data and null values are not a huge problem in isolation, but a growing number of empty values may help determine whether to drop a variable, ignore missing values, or compute a predicted value. Most data scientists support the dropping of variables if an attribute has upward of 50% empty or null values. This threshold does vary.

Low Variance

Data attributes that are very similar to one another do not carry much information. As a result, a lower variance threshold can be set and dimensions removed based on this. As variance is dependent on range, data normalization should take place first.

High Correlation

Attributes with similar data trends are likely to carry similar information. Multicollinearity is the carrying of similar information that can reduce the performance of the model. Correlation between continuous data is determined by the Pearson's product-moment coefficient, which is a measure of the linear correlation between two variables X and Y. For discrete data, the Pearson's chi-square value determines how likely it is that any observed difference between sets arose by chance.

Random Forest Decision Trees

The use of random forest decision tree ensembles is useful for identifying key features. Random forest decision trees use a subset of features that can identify best split attributes. If an attribute is often selected as the best split node, it is likely to be a feature to keep. Decision trees are particularly useful for visualizing reductions also.

Backward Feature Elimination

In Backward Feature Elimination, the model trains on n attributes. Starting with n attributes, at each iteration, the model is trained on $n-1$ attributes n number of times. The attribute that demonstrates the smallest increase in error rate is removed, and the algorithm performs another iteration on $n-1$ attributes. At each iteration k , a model with $n-k$ attributes is trained. This is computationally expensive.

Forward Feature Construction

Forward Feature Construction is the opposite of Backward Feature Elimination. The model starts with one attribute and evaluates which of the attributes has the highest increase in performance. Much like its opposite Backward Feature Elimination, this method is computationally expensive unless operating in a lower range of dimensions.

Principal Component Analysis (PCA)

Principal component analysis decreases the number of variables (axes) through transforming an original set of variables into a new set, or the principle components (see Figure 4-19). This process is iteratively calculating the maximum variance within features in the data. Each

component is orthogonal and a linear combination of original variables. Orthogonality indicates that the correlation between the components is zero.

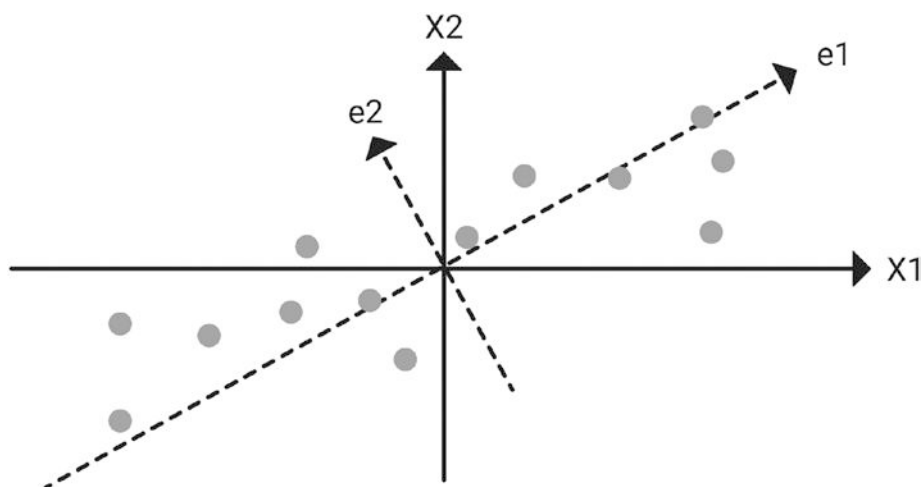


Figure 4-19. PCA

The first principle component is the linear combination of the original dimension that has the maximum variance; and the n th principal component is the linear combination with the highest variance, subject to being orthogonal to the $n - 1$ principle components.

Natural Language Processing (NLP)

Natural Language Processing or NLP pursues a defined set of problems within AI. NLP is defined as the ability of systems that can analyze, understand, and generate human language, including speech and text.

NLP is an aspect of computational linguistics (studying linguistics using computer science) and is useful in the following:

- The retrieval of structured and unstructured data within a dataset. For example, searching clinical notes by keyword or phrase
- Social media monitoring
- Question answering: interpretation of natural language from humans to interact appropriately; for instance, as with virtual assistants or speech recognition software
- Analysis of a document to determine key findings
- Ability to parse and interpret a text to understand sentiment and mood
- Recognizing distinctions among diagnoses and relationships
- Image to text recognition; for instance, reading a sign or menu
- Machine translation: NLP is used in machine translation programs in which one human language is automatically translated into another human language
- Topic modeling—What is this document talking about?
- Understanding sentiment from social media or discussion posts

Interpreting natural language is fraught with challenges, as human language is naturally ambiguous—language, pronunciation, expression, and perception. Although there are rules with human language, they are often misunderstood and misused. NLP takes into consideration

the structure of language to derive meaning. Words make phrases; phrases make sentences; sentences make documents; and all of the aforementioned convey ideas.

NLP has a tool kit of text processing procedures including a range of data mining methods that can be used for model development. Due to the nature of unstructured data, NLP tasks can be expensive regarding computational resource and time. Neural networks and deep learning can also be used for NLP tasks.

With most data generated existing in the form of unstructured data, NLP is a powerful tool to interpret and understand natural language.

As with any aspect of computing, there are several terms to understand before proceeding:

- **Tokenization:** the process of converting a corpus of text into smaller units, or tokens; there are many algorithms available for breaking a text into tokens
- **Tokens:** words or entities present in the text
- **Text object:** a sentence or a phrase or a word or an article
- **Stemming:** a basic rule-based process of stripping suffixes (“ing,” “ly,” “es,” “s,” etc.) from words
- **Stem:** the text created after stemming
- **Lemmatization:** determining the root of a word from dictionaries and morphological analysis
- **Morpheme:** unit of meaning in a language
- **Syntax:** arranging symbols (words) to make a sentence; it involves determining the structural role of words in the sentence and phrases

- Semantics: meaning of words and how to join words into meaningful phrases and sentences
- Pragmatics: using and understanding sentences in different situations and how interpretations are affected

In this section, I aim to introduce key concepts and methods of NLP and to demonstrate the techniques that can be applied to datasets to generate defined value.

Getting Started with NLP

Giving an NLP model an input sentence and receiving a useful output requires several key components (see Figure 4-20).

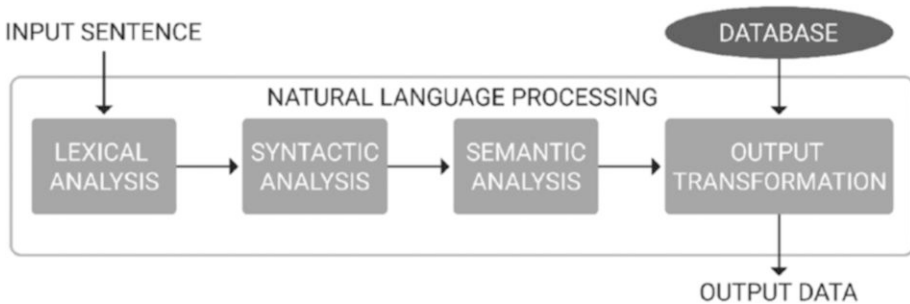


Figure 4-20. *How NLP works*

Preprocessing: Lexical Analysis

As with any dataset, a corpus of text that is not relevant to the context of data can be understood as noise. The first stage in NLP is to clean and standardize the input text, ensuring it is noise free and ready for analysis.

Over and above spelling correction and grammar correction, the following techniques are used to reduce noise.

Noise Removal

Noise removal involves preparing a dictionary of noisy tokens (i.e., words) and parsing the text, removing the tokens found in the noise dictionary. For instance, words like the, a, of, this, that, and so forth, would be removed.

Lexicon Normalization

Follow, following, followed, follower are all variations of the word follow. Contextually, the words are similar. Lexicon normalization reduces dimensionality through stemming, which strips suffixes and prefixes, and lemmatization, which is a defined procedure that uses word structure and grammar relationships.

Porter Stemmer

Python: NLTK; Method: PorterStemmer

The Porter stemmer algorithm is a popular and useful method to improve the effectiveness of information retrieval.[54] The algorithm works on the principle that many words in English share a common root. Stemming works on suffixes and removes common morphological and inflexional endings from words. Therefore, stemming allows one to reduce similar words into a common root form.

For example, take the text, “I felt troubled by the fact that my best friend was in trouble. Not only that, but the issues I had dealt with yesterday were still troubling me.”

The words troubled, trouble, and troubling all share the common root trouble. Therefore, according to the Porter stemmer algorithm, instead of counting all three words once, the stem trouble is counted thrice instead. The benefit of stemming is that common words can be clustered under

a common stem to provide a more accurate statistical representation of the number of occurrences of a certain word. However, a drawback of stemming is that the semantic meaning of the word may be lost.

Stemming and lemmatization are used to reduce inflectional forms and sometimes derivationally related forms of a word to a common base form.

Object Standardization

A corpus of text may contain words that cannot be found in lexicon dictionaries. For example, on Twitter, someone may mention DM'ing someone, or another may like an RT of someone else's tweet. Acronyms, hashtags, slang, and colloquialisms can be removed through prepared dictionaries or using regular expressions.

Syntactic Analysis

To be analyzed, the text needs to be converted into features. The syntactic analysis of text involves analyzing sentences to understand relationships between words and assigning a syntactic structure to it. There are several algorithms for syntactic analysis; however, the Context-Free Grammar is most popular due to it being the simplest style of grammar and therefore widely used.

Take the sentence, "David saw a patient with uncontrolled type 2 diabetes."

Within the sentence, we need to identify the subject, objects, noise, and attributes to understand the sequence of words and its dependencies.

Dependency Parsing

Toolkit: NLTK; Method: StanfordDependencyParser

Sentences are composed of words in a structure. Basic grammar can determine the relationship or dependencies between a structure. Dependency parsing represents this in a tree structure and represents grammar and arrangement of words (see Figure 4-21).

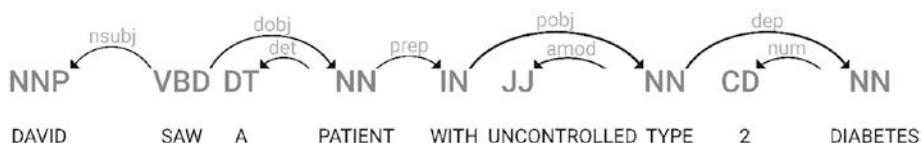


Figure 4-21. *Dependency parsing*

Dependency grammar analyses asymmetrical binary relationships between tokens. The Stanford parser from NLTK is commonly used for this purpose.

Using the example in Figure 4-21, the parse tree determines the root of the word as “saw” and is then linked by subtrees. The subtrees are split by subject and object, with each subtree also showing dependencies.

Part of Speech Tagging

Toolkit: NLTK; Method: word_tokenize

Part of speech tagging involves associating each word or token in a sentence with a part of speech (POS) tag. These tags are the basic English labels that you learned in primary school and determine nouns, verbs, adjectives, adverbs, numbers, and so on.

This is particularly useful for tasks such as building parse trees—which can, in turn, be used for determining what something is, sentiment analysis, determining appropriate answers to questions, or understanding similar entities (see Figure 4-22).

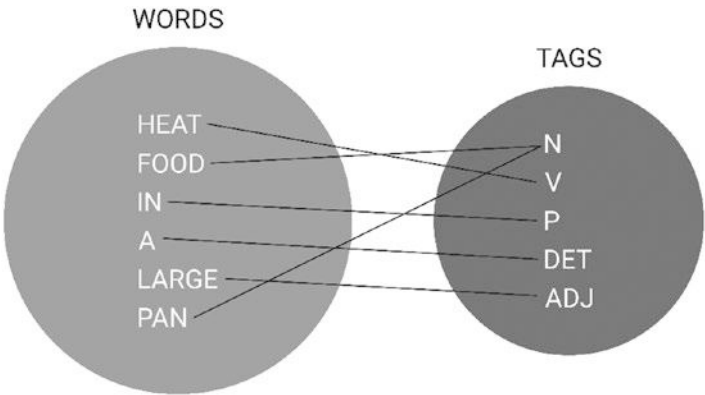


Figure 4-22. *Part of speech tagging*

Part of speech tagging is the basis of the assists in several areas.

Reducing Ambiguity

Some sentences have multiple meanings given the structure; for example, take the two sentences:

“I managed to read my book on the train.” “Can you please book my train tickets?”

Part of speech tagging identifies “book” as a noun in the first sentence and as a verb in the second sentence.

Identifying Features

By identifying the types of speech alongside different contexts of a word, POS can distinguish between uses and creates stronger features for use.

Normalization

POS tags are the foundation of normalization and lemmatization, to understand sentence structure and dependency.

Stopword Removal

POS is useful in removing commonly used words, or stopwords, from a text.

Semantic analysis

Semantic analysis is the most complex phase of NLP. It draws the exact meaning or the dictionary meaning from the text. Using knowledge about the structure of words and sentences among the context, the meaning of words, phrases, sentences, and texts is stipulated, and subsequently, also their purpose and consequences.

Techniques Used Within NLP

Once data preprocessing, lexical, syntactical, and semantic analysis of corpus has taken place, we are required to transform text into mathematical representations for evaluation, comparison, and retrieval. For instance, searching a collection of patient profiles for users with “hypertension” should only bring out those with hypertension. This is achieved through transforming documents into the vector space model with scoring and term weighting essential for query ranking and search retrieval. Documents can be in the form of patient records, web pages, digitalized books, and so forth. The following algorithms are typical in the comparison and evaluation process.

N-grams

Python: NLTK; Method: ngrams

N-grams are used in many NLP problems. If X = number of words in a given sentence K , the number of n -grams for sentence K would be

$$N \text{ grams}_K = X - (N - 1)$$

For example, if $N = 2$ (bigrams), the sentence “David reversed his metabolic syndrome” would result in n-grams of the following:

- David reversed
- Reversed his
- His metabolic
- Metabolic syndrome

N-grams preserve the sequences of N items from the text input. N-grams can have a different N value: unigrams for when $N = 1$, bigrams when $N = 2$, and trigrams when $N = 3$.

N-grams are used extensively for spelling correction, word decomposition, and summing texts.

TF IDF Vectors

A document can be represented as a high-dimensional vector in the space of words. Each entry in the vector corresponds to a different term within the documents and the number of its occurrences. The TF-IDF (Term Frequency-Inverse Document Frequency) vector scheme assigns each term within the document a weight using the following formula:

$$W_{td} = f_{td} \cdot \log \left(\frac{D}{ND_t} \right)$$

Each term’s weight (W_{td}) is calculated by multiplying the frequency of the term (f_{td}) by the log of the total number of documents (D) divided by the number of documents the term occurs at least once in (ND_t). The ordering of the terms may not necessarily be maintained. Weights for the terms gathered can then be used to determine documents with high frequencies of the specific term within them. A collection of TF-IDF vectors could be used to represent a user’s interest.

Weights for the terms gathered can then be used to determine documents with high frequencies of the specific term within them.

Latent Semantic Analysis

Where there is a corpus of a large number of documents, for each document d , the dimension of the vector representing each document can typically exceed several thousand. Latent semantic analysis relies on the fact that intuitively, terms in documents may often be related. For example, if document d contains the term *sea*, it will often contain the word *beach*.

Equivalently, if the vector representing d has a non-zero component in the entry for *sea*, it will also have a non-zero component in the *beach* component. If this kind of structure can be detected, relationships between words can be automatically learned from the data.

The word-document is represented by a matrix A , which is decomposed using Singular Value Decomposition to give the strength of the most significant correlations and their directions.

The decomposition of A allows one to discover the semantics of the document through correlations between terms and their significance within the document. The latent semantic analysis method can be applied to determine the context of a variety of materials to present a web user with results that were contextually relevant.

Cosine Similarity

Python: SciPy; Method: cosine

The cosine similarity is a metric that measures the similarity between two vectors. Therefore, this could either be used to define the similarity between two documents, or a document and a query. The cosine similarity between two vectors can be defined as the following:

$$Sim(u,v) = \frac{u \cdot v}{\|u\| \cdot \|v\|}$$

The similarity between two vectors is calculated as the inner product of the vectors, divided by the product of the lengths of the vectors in question. Intuitively, the greater the angle between two documents, the less similar they are. This holds for vectors in any N-dimensional space. Furthermore, the TF-IDF vector scheme can be integrated with the cosine similarity metric to determine the similarity between a search query and a plethora of documents:

$$Sim(q,d) = \frac{\sum_{t \in q \cap d} W_{td} \cdot W_{tq}}{\|d\| \cdot \|q\|}$$

The similarity between a query q and document d is calculated as the product of the TF-IDF weights for the term in both the document and query summed over all the terms in both the query and document; this is divided by the product of the length of the document and length of the query. Practically, however, calculating $Sim(q,d)$ would prove computationally expensive as the number of documents used grows.

Naïve Bayesian Classifier

The Bayesian classifier is based on the Bayesian theorem and is particularly suited when the dimensionality of the inputs is high. Despite its simplicity, Naïve Bayes can often outperform more sophisticated classification methods. Given a specified threshold, this method can be used to classify the probability of whether a vector representing a document is of interest to a user. Given an attribute d , we can calculate whether the example belongs in class C with the following formula:

$$P(C=c|D=d) = \operatorname{argmax}_c \frac{P(D=d|C=c)P(C=c)}{P(D=d)}$$

Other techniques such as kNN and ANN can also be used to classify and retrieve information.

Genetic Algorithms

Genetic algorithms (GA) are a fascinating topic within machine learning. GA take inspiration from evolution to minimize the error rate, attempting to mimic the function of chromosomes much like neural networks attempt to mimic the human brain.

Evolution is considered the optimal learning algorithm. In machine learning, the application of this is in models whereby several candidate answers (referred to as chromosomes or genotypes) are produced, and the cost function applied to all. In GA, a fitness function is defined that determines if the chromosomes fit enough to mate. Chromosomes furthest away from the optimal outcome are removed. Chromosomes are also subject to mutation. GA are a type of search and optimization learner and apply to discrete and continuous problems.

Chromosomes that are close to the optimal solution may be combined. The combination or mating of chromosomes is known as a crossover. The survival of the fittest approach identifies chromosomes that aim to express characteristics that adhere to natural selection—where the offspring is more optimal than the parent.

Mutation helps to overcome overfitting. It is a random process to get over local optima and find the global optimum. Mutation helps to ensure that child chromosomes are different from the parents' and continue evolution.

The degree in which chromosomes mutate and mate are parameters that can be controlled or left to the model to learn.

GA have varied applications:

- Detection of blood vessels is ophthalmology imaging
- Detecting the structure of RNA
- Financial modeling
- Routing vehicles

A group of chromosomes is referred to as a population. Although it stays at a defined, constant size, it usually evolves to better average predictions over the course of generations or time.

The evaluation of a chromosome, c , is calculated as its evaluation function value divided by the average of the generation, represented as the following:

$$\text{fitness}(c) = g(c) / (\text{average of } g \text{ over the entire population})$$

John Holland invented the GA approach (see Figure 4-23) in the early 1970s.[55] The automation of chromosome selection is known as genetic programming.

Best Practices and Considerations

Mastering the art of machine learning takes time and experience. However, there are several areas worthy of consideration, particularly for beginners to machine learning, to ensure the best use of time and optimal model efficiency.

Good Data Management

With the volume and complexity of data used in machine learning tasks, proper data management is paramount. Not only does this involve having the processes, policies, procedures, permissions, and certifications

required for data management, it consists of the managing and transformation of training data to learn an optimal model.

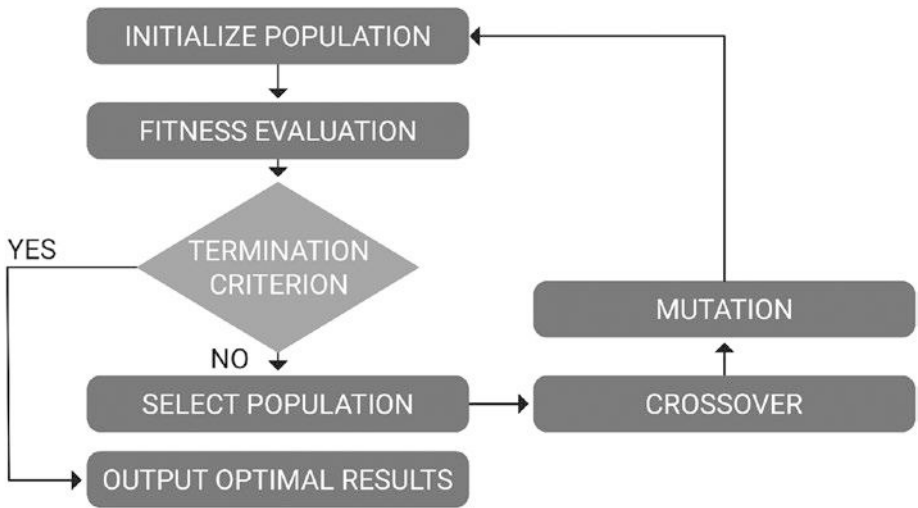


Figure 4-23. Basic structure of a genetic algorithm

Establish a Performance Baseline

It is crucial to have a baseline to measure the performance of your algorithm against other iterations and model types. As there is no one perfect algorithm for every problem, attempt many algorithms to identify the relative performance of your models.

Spend Time Cleaning Your Data

The time required to train a machine learning model varies significantly between algorithms. The accuracy of data and time spent on training the model positively influence model accuracy. Take time to clean your data to ensure predictions are as robust as possible.

It is often the case that not all of the input variables impact the dependent variable or outcome. Ensure the variables used by the model do not include those that are irrelevant.

Training Time

If the dimensionality of data is large and computing power limited, training time will be extensive. When time is short, it is useful to consider machine learning algorithms that do not require substantial training. For example, neural networks may not be appropriate for time-limited tasks.

Choosing an Appropriate Model

Some machine learning algorithms make particular assumptions about the structure of data or the desired outputs. For this reason, it is helpful to consider the model choice and whether it is an appropriate approach. Utilizing the right model has a plethora of benefits, including more accurate predictions, faster training times, and more useful results.

Some machine learning models are resistant to outliers or nonparametric tests. For example, decision-tree-based approaches typically classify a node into two based on a threshold. Data outliers are therefore less impactful in tree-based approaches.

Evaluating a range of models is a useful approach to machine learning. Occam's Razor is typically applied to choosing the model of choice. Occam's Razor seeks to identify the most straightforward model that achieves the desired output given all else being equal.

Choosing Appropriate Variables

Although more data is ordinarily invited in machine learning problems, it is typically preferable to work with fewer predictor variables for many reasons.

Redundancy

Increasing the number of variables within a training dataset increases the chances of models learning hidden relationships between them. It is vital to identify unnecessary variables and only use nonredundant predictor variables within models. A model that is learning redundant connections will impact model accuracy.

Overfitting

Even if there are no relationships among predictor variables within a model, it is still beneficial to use fewer variables. Complex models, or those that use a high number of predictor variables, typically suffer from overfitting. As a result, models perform well on training datasets but are less accurate on validation and real-world settings, as they learn the error within the data (i.e., the noise) rather than the signal, or relationships, between variables.

Productivity

Even if all of the variables within a sophisticated machine learning model are relevant, there is a practical impact in using a large number of predictor variables that can affect productivity. Practical considerations include the amount of data available, the subsequent effect on storage, computing resources, associated costs, time allocated for the project, and the time required for learning and validation. Predictor variables can be identified through feature selection, or transformed through feature extraction. SVMs are useful in cases of increased data dimensionality.

The Pareto principle (or the 80/20 rule or law of a vital few) is a useful measure to use in machine learning projects. The Pareto principle states that approximately 80% of the effects come from 20% of the causes. Using

the Pareto principle, focusing on the 20% most significant predictor variables should facilitate the building of relatively successful models within a reasonable time.

Understandability

Models with fewer predictor variables are easier to visualize, understand, and explain. A key aspect of a successful machine learning project is that all stakeholders can comprehend the model. This often requires data scientists trading off. As the result of reducing the number of predictor variables, there is a chance there may be a reduction in the success of a machine learning model. However, this concurrently makes the model easier to interpret and understand. The usefulness of this approach is only realized toward the end of a project when it comes to sharing not only the performance of a model but also how it works. This is particularly relevant in healthcare, as there are concerns about using black-box models.

Accuracy

Any machine learning model aims to generalize well. Depending on the use case, an approximation may be more beneficial than a precise, accurate output. Through approximating values, models tend to avoid overfitting and reduce the amount of time spent on processing.

Impact of False Negatives

When evaluating the impact of a model before deployment into a live environment, consider the impact of false negatives. For example, take a predictive model that classifies the risk of breast cancer. A false positive would mean that a patient is informed they have breast cancer when

they do not, which should be identified later in the treatment pathway. However, a false negative would mean a patient with breast cancer would not be notified—which is potentially far worse and costly.

Linearity

Many machine learning algorithms assume that relationships are linear: in other words, that classes can be separated by a straight line of best fit or its higher dimensional representation. There are instances where relationships are not linear. Relying on a linear classification algorithm where there is a nonlinear class boundary will result in low accuracy.

Linear regression, for example, assumes both input and output variables do not contain noise. It is vital to expose the signal within data to ensure models learn the signal or the correct relationships within the data. It is advantageous to identify and remove outliers for the output variable in particular to reduce the chances of learning the noise.

Data with a nonlinear trend may require transforming to make the relationship linear: for instance, log-transforming data where there is an exponential relationship.

Linear algorithms are typically the first methods attempted in machine learning scenarios. As a result of linearity, algorithms are simple and quick to train.

Parameters

Each machine learning model is subject to parameters and hyperparameters. The time and resource required to train a model increases as the number of parameters does. An algorithm's parameter is a variable that is utilized by the model where the value can be estimated from the dataset.

An algorithm's hyperparameter is a variable that is external to the model and used to estimate parameters. A hyperparameter's value cannot be estimated from data. Instead, it is set by the data scientist, randomly or otherwise, or through the use of heuristics.

Ensembles

In specific machine learning problems, it may prove more useful to group classifiers together using the techniques of voting, weighting, and combination to identify the most accurate classifier possible. Ensemble learners are very useful in this aspect.

Use Case: Type 2 Diabetes

Type 2 diabetes is one of the most significant health and economic burdens facing the global population, with 1 in 6 people dying from diabetes or its related complications every 6 seconds.[56]

There are many factors involved in the progression of the disease, which, if tackled, can help to prevent its progression through targeted treatment profiling, leading to lower patient morbidity and mortality rate. Data in the form of health biomarkers—typically blood glucose, HbA1c, fasting blood glucose, insulin sensitivity, and ketones—are used to understand and monitor disease burden and response to treatments.

As a result of the tremendous burden of type 2 diabetes, there has been a substantial investment in developing intelligent models within this area, with much of the focus on diagnosing, predicting, and managing the condition through machine learning and data mining.

Many machine learning techniques have been applied to type 2 diabetes datasets. Traditional techniques, ensemble techniques, and unsupervised learning—in particular, association rule mining—has been used to identify prediction models with optimal accuracy.

Farideh Bagherzadeh-Khiabani et al. used the clinical dataset from 803 female prediabetics consisting of 55 features collected over a decade to develop multiple models for predicting type 2 diabetes risk.[57] The study developed a logistic model, demonstrating wrapper models, or models that consider a subset of features, improve the performance of clinical prediction models. This was developed into an R program to visualize output.

Similarly, Georgia El et al. used random forest decision trees on many features collected from 15 people with type 1 diabetes to predict short-term subcutaneous glucose concentrations.[58] Glucose concentration was predicted through the use of SVMs. The study concluded that the inclusion of two biomarker features—8-hydroxy-2-deoxyguanosine, an oxidative stress marker, and interleukin-6—improved classification accuracy.

Duygu Çalışır et al. developed an automatic diagnosis system for type 2 diabetes with a classification accuracy of almost 90% through the use of linear discriminant analysis and SVM classifiers.[59] Latent Dirichlet Allocation (LDA) was used to distinguish features between healthy patients and patients with type 2 diabetes, which was then in turn used as the input for the SVM classifier. A tertiary stage evaluated sensitivity, classification, and confusion before determining the output.

Razavian et al. conducted a more extensive study, determining that type 2 diabetes prediction models could be accurately derived from population-scale datasets.[60] Over 42,000 variables were collected from over 4 million patients between 2005 and 2009. Machine learning was used to determine relevant features, which tailed out at around 900. Novel risk factors for type 2 diabetes were identified, such as chronic liver disease, high alanine aminotransferase, esophageal reflux, and history of acute bronchitis.

Machine learning has likewise been applied to diagnosing the risk of other diseases related to diabetes. Lagani et al. identified the least set of clinical variables with the best predictive accuracy for complications such as cardiovascular disease (heart disease or stroke), hypoglycemia, ketoacidosis, proteinuria, neuropathy, and retinopathy.[61]

Huang et al. used decision-tree-based prediction models on identifying diabetic nephropathy in patients.[62] Leung et al. compared several methods: partial least square regression, regression trees, random forest decision trees, naïve Bayes, neural networks, and SVMs on genetic and clinical features.[63] Patient age, years of diagnosis, blood pressure, genetic polymorphisms of uteroglobin, and lipid metabolism arose as the most efficient predictors of type 2 diabetes.

Summers et al. developed a predictive model utilizing a dataset consisting of over 120 features to identify comorbid depression and diabetes-related distress. Summers et al. explored several techniques including random forest decision trees, SVMs, naïve Bayes, and neural networks. Patient age and employment status were defined as the variables most likely to affect predictive accuracy.[64]

Unsupervised learning has mainly concentrated on association rule mining, identifying associations between risk factors in people with diabetes. Simon et al. proposed an extension to association rule mining labeled Survival Association Rule Mining, which included survival outcomes, and made adjustments for confounding variables and dosage effects.

CHAPTER 5

Evaluating Learning for Intelligence

“Intelligence is the ability to adapt to change”

—Stephen Hawking

As a field, machine learning is still in its infancy. Advanced machine learning has only been explored over the last 25 years, which has fueled data science as a profession. As a result, the data science industry is still in a phase of wonderment at the endless potential of AI and machine learning. With this comes both excitement and confusion—and an industry that is gathering knowledge, experience, and first-time problems.

Typically, the most laborious tasks within a machine learning project are identifying the appropriate model and engineering features, which make a substantial difference to the output of the model. In fact, the features chosen can often have more impact on the quality of a model compared to the model choice itself. Therefore, it is important to evaluate the learning algorithm that will determine the model’s intelligence to predict the output of an unknown sample. This is usually done using various metrics, which are discussed further.

Model Development and Workflow

To successfully deploy a machine learning model, there are several stages of development and evaluation that take place, as illustrated in Figure 5-1.

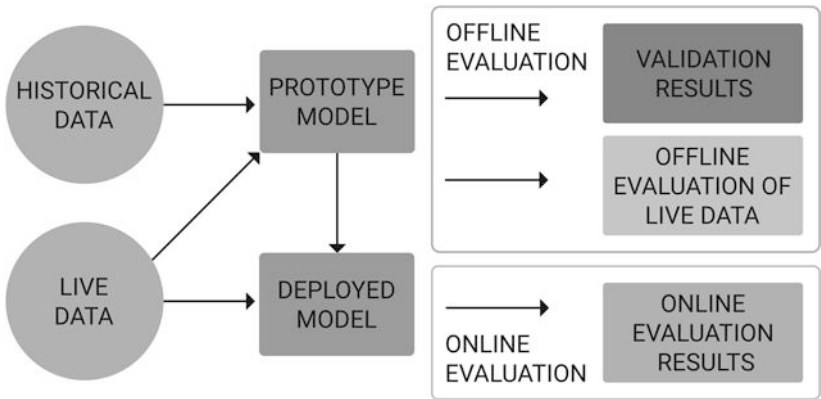


Figure 5-1. *Model development and work flow*

The first stage is the prototype phase. During this phase, a prototype is created through testing various models on historical data to determine the best model. Hyperparameter tuning, as discussed later in this chapter, is a requirement of model training. Once the best prototype model is chosen, the model is tested and validated. Validating a model requires splitting datasets into training, testing, and validation sets as discussed in Chapter 3. Consider the fact that there is no such thing as a random dataset and instead the randomness applies to the splitting of the dataset. Be aware of biases that may appear in the data. Once the model has been successfully validated, it is deployed to production. The model is then usually evaluated by one (or several) performance metrics.

There are two ways of evaluating a machine learning model: offline evaluation and online (or live) evaluation.

Why Are There Two Approaches to Evaluating a Model?

A deployed machine learning model consumes data from two sources: historical data (or the data that is used as the experience to be learned from) and live data. Many machine learning models assume stationary distribution data—that the data distribution is constant over time.

However, this is atypical of real life, as distributions of data often change over time—known as a distribution shift. For instance, consider a system that predicts the side effects of medications to patients based on their health profile. Medication side effects may change based on population factors such as ethnicity, disease profile, territory, medication popularity, and new medications. The distribution of relevant side effects based on patient data can vary quickly over time, and hence it is essential for a model to detect a shift in distribution and accordingly evolve the model. The method in which this is typically assessed is through the performance of the model based on live data, evaluated through the validation metric used in the testing and validation of the model on historical data.

Model performance that is similar to or within a threshold of permissibility when evaluated on live data is deemed as a model that continues to fit the data. Degradation of model performance indicates that the model does not fit the data and requires retraining.

Offline evaluation measures the model based on metrics learned and evaluated from the historical, stationary, distributed dataset. Metrics such as accuracy and precision-recall are typically used within the offline training stage. Offline evaluation techniques include the hold-back method and n-fold cross-validation.

Online evaluation refers to the evaluation of metrics once the model is deployed. The key takeaway is that these metrics may differ from the metrics used to evaluate performance when the model is deployed live. For instance, a model that is learning on new pharmacological treatments may seek to be as precise as possible in training and validation; but when placed online, it may need to consider business goals such as budget or treatment value when deployed. Online evaluation, particularly in the digital age, can support multivariate testing to understand best-performing models.

Feedback loops are key to ensuring systems are performing as intended and help to understand the model in the context of use better. This can be performed by a human agent or automated through a contextually intelligent agent or users of the model.

It is important that the evaluation of a machine learning model is based on a statistically independent dataset and not on the dataset it is trained on. This is because the evaluation of the training dataset is optimistic about the model's true performance as it adapts to the dataset. By evaluating the model with previously unseen data, there is a better estimate of the generalization error. New data can be hard to find; hence it is important to be able to have new, unseen data from the current dataset. Methods such as n-fold cross-validation discussed in Chapter 3 are useful techniques for this purpose. Often the data used is more important than the algorithm choice; and the better the features used, the greater the performance of the model.

The evaluation metrics discussed can be found in the metrics package for R and scikit-learn for Python.

Evaluation Metrics

There are a plethora of evaluation metrics for machine learning problems. Metrics exist for the variety of machine learning tasks—classification, regression, clustering, association rule mining, NLP, and so on.

Classification

Classification problems seek to give a label or classification to an input. There are several methods by which to measure performance, including accuracy, precision-recall, confusion matrices, log-loss (logarithmic loss), and AUC (area under the curve).

Accuracy

Accuracy is the simplest technique used in identifying whether a model is making correct predictions. It is calculated as a percentage of correct prediction over the total predictions made.

Accuracy = number of correct predictions/number of total predictions

Confusion matrix

Accuracy is a general metric that does not consider the division between classes. Therefore, it does not consider misclassification or the associated penalty with misclassification. For instance, a medical misdiagnosis that is a false positive (e.g., take a patient diagnosed with breast cancer when they do not have it) has substantially different consequences compared to a false negative, whereby a patient is told that they do not have breast cancer when in fact they do. A confusion matrix breaks down the correct and incorrect classifications made by the model and attributes them to the appropriate label.

- True positive: Where the actual class is yes, and the value of the predicted class is also yes.
- False positive: Actual class is no, and predicted class is yes

- True negative: The value of the actual class is no, and the value of the predicted class is no
- False negative: When the actual class value is yes, but predicted class is no

Take an example whereby a model predicts whether a patient has breast cancer or not based on 50 example inputs from the test dataset with an equal distribution between positive and negative labeled examples. The confusion matrix would be as in Table 5-1.

Table 5-1. *Confusion matrix*

	Prediction: Positive	Prediction: Negative
Labeled positive	20	5
Labeled negative	15	10

From the confusion matrix, it is determined that the positive class has greater accuracy than the negative class. The accuracy of the positive classification is $20/25 = 80\%$. The negative class has an accuracy of $10/25 = 40\%$. Both metrics differ from the overall accuracy of the model, which would be determined as $(20 + 10)/50 = 60\%$. It is apparent how a confusion matrix adds more detail to the overall accuracy of a machine learning model.

As a result, accuracy can be rewritten as the following:

Accuracy = (correctly predicted observation)/(total observation) = $(TP + TN)/(TP + TN + FP + FN)$

Per-class accuracy

Per-class accuracy is an extension of accuracy that takes into account the accuracy of each class. As a result, the preceding example has a per-class accuracy of $(80\% + 40\%)/2 = 60\%$. Per-class accuracy is useful

in distorted problems where there are a larger number of examples within one particular class compared to another. The class with greater examples dominates the calculation, and therefore accuracy alone may not suffice for the nature of your model; thus it is useful to evaluate per-class accuracy also.

Logarithmic loss

Logarithmic loss (or log-loss for short) is used for problems where a continuous probability is predicted rather than a class label. Log-loss provides a probabilistic measure of the confidence of the accuracy and considers the entropy between the distribution of true labels and predictions.

For a binary classification problem, the logarithmic loss would be calculated as follows:

$$\text{Log-loss} = -\frac{1}{N} \sum_{i=1}^N y_i \log p_i + (1 - y_i) \log(1 - p_i)$$

Where P_i is the probability of the i th data point belonging to a class and y_i the true label (either 0 or 1).

Area Under the Curve (AUC)

The AUC plots the rate of true positives to the rate of false positives. The AUC enables the visualization of the sensitivity and specificity of the classifier. It highlights how many correct positive classifications can be gained allowing for false positives.

The curve is known as the receiver operating characteristic curve, or ROC as shown in Figure 5-2. A high AUC or greater space underneath the curve is good, and a smaller area under the curve (or less space under the curve) is undesirable. In Figure 5-2, test A has better AUC as compared to test B, as the AUC for test A is larger than for test B. The ROC visualizes the trade-off between specificity and sensitivity of the model.

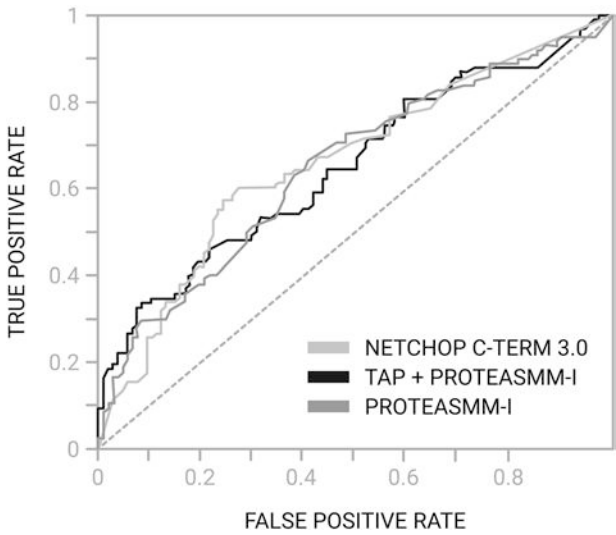


Figure 5-2. ROC curve

Precision, recall, specificity, and F-measure

Precision and recall are two metrics used together to evaluate model performance. Precision evaluates how many items are truly relevant compared to the total number of items correctly classified. Recall evaluates how many items are predicted to be relevant by the model from the items that are relevant.

- Precision: (correctly predicted Positive)/(total predicted Positive) = $TP / (TP + FP)$
- Recall: (correctly predicted Positive)/(total correct Positive observation) = $TP / (TP + FN)$

Specificity refers to how well the model performs at returning incorrect classifications and is calculated as in Figure 5-3.

- Specificity: (correctly predicted Negative)/(total Negative observation) = $TN / (TN + FP)$

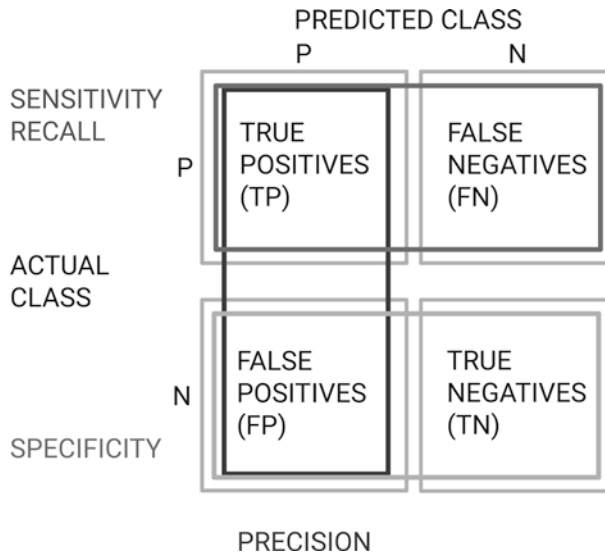


Figure 5-3. Specificity classification diagram

F-measure goes beyond the arithmetic mean and calculates the harmonic mean of precision and recall:

$$F = \frac{1}{\frac{1}{2} \left(\frac{1}{p} + \frac{1}{r} \right)} = \frac{2pr}{p+r}$$

Where p denotes precision and r denotes recall.

Regression

Regression machine learning models output continuous variables, and root-mean-squared error (RMSE) is the most commonly used evaluation metric for these problems.

RMSE

RMSE calculates the square root of the sum of the average distance between predicted and actual values. This can also be understood as the average Euclidean distance between the true value and predicted value vectors. A criticism of RMSE is that it is sensitive to outliers.

$$RMSE = \sqrt{\frac{\sum_i (y_i - \hat{y}_i)^2}{2}},$$

where y_i denotes the actual value and \hat{y}_i denotes predicted value.

Percentiles of errors

Percentiles (or quantiles) of error are more robust as a result of being less sensitive to outliers. Real-world data is likely to contain outliers, and thus it is often useful to look at the median absolute percentage error (MAPE) rather than the mean.

$$MAPE = median\left(\left\| \left(y_i - \hat{y}_i \right) / \left(y_i \right) \right\| \right),$$

where y_i denotes the actual value and \hat{y}_i denotes predicted value.

The MAPE is less affected by outliers by using the median of the dataset. A threshold or percentage difference for predictions can be set for a given problem to give an understanding of the precision of the regression estimate. The threshold depends on the nature of the problem.

Skewed Datasets, Anomalies, and Rare Data

An experienced data scientist treats all data with suspicion. Data can be inconsistent; and as a result, skewed datasets, imbalanced class examples, and outliers can all significantly affect the performance of a model. Having more examples within one class compared to another can lead to an underperforming model. Furthermore, outliers or data anomalies can further skew performance evaluation metrics. The effect of large outliers can be mitigated using percentiles of error. In practice, good data cleansing, removal of outliers, and normalization of variables can reduce the sensitivity to outliers.

Parameters and Hyperparameters

Hyperparameters and parameters are often used interchangeably, yet there is a difference between the two. Machine learning models can be understood as mathematical models that represent the relationship between aspects of data.

Model parameters are properties of the training dataset that are learned and adjusted during training by the machine learning model. Model parameters differ for each model, dataset properties, and the task at hand. For instance, in the case of an NLP predictor that output the sophistication of a corpus of text, parameters such as word frequency, sentence length, and noun or verb distribution per sentence would be considered model parameters.

Model hyperparameters are parameters to the model building process that are not learned during training. Hyperparameters can make a substantial difference to the performance of a machine learning model. Hyperparameters define the model architecture and effect the capacity of the model, influencing model flexibility.

Hyperparameters can also be provided to loss optimization algorithms during the training process. Optimal setting of hyperparameters can have a significant effect on predictions and help prevent a model from overfitting. Optimal hyperparameters often differ between datasets and models.

In the case of a neural network, for example, hyperparameters would include the number and size of hidden layers, weighting, learning rate, and so forth. Decision trees hyperparameters would include the desired depth and number of leaves in the tree. Hyperparameters with a support vector machine would include a misclassification penalty term.

Tuning Hyperparameters

Hyperparameter tuning or optimization is the task of selecting a set of optimal hyperparameters for a machine learning model. Optimized hyperparameters values maximize a model's predictive accuracy. Hyperparameters are optimized through running training a model, assessing the aggregate accuracy, and appropriately adjusting the hyperparameters. Through trialing a variety of hyperparameter values, the best hyperparameters for the problem are determined, which improves overall model accuracy.

Hyperparameter Tuning Algorithms

Hyperparameter tuning is like training a machine learning model. The task at hand is one of optimization. Model parameters can be expressed as a loss function, whereas hyperparameters cannot be expressed as such, as it depends entirely on the model training process. There are several approaches to hyperparameter tuning, with the most common being grid search and random search.

Grid Search

The grid search is a simple, effective, yet resource expensive hyperparameter optimization technique that evaluates a grid of hyperparameters. The method evaluates each hyperparameter and determines the winner. For example, if the hyperparameter were the number of leaves in a decision tree, which could be anywhere from $n = 2$ to 100, grid search would evaluate each value of n (i.e., points on the grid) to determine the most effective hyperparameter.

It is often a case of guessing where to start with hyperparameters, including minimum and maximum values. The approach is typical of trial and error, whereby if the optimal value lies toward either maximum or minimum, the grid would be expanded in the appropriate direction in an attempt to further optimize the model's hyperparameters.

Random Search

Random search is a variant of grid search that evaluates a random sample of grid points. Computationally, this is far less expensive than a standard grid search. Although at first glance it would appear that this is not as useful in finding optimal hyperparameters, Bergstra et al. demonstrated that in a surprising number of instances, a random search performed roughly as well as grid search.[65] The simplicity and better-than-expected performance of a random search means that it is often chosen over grid search. Both grid search and random search are parallelizable.

More intelligent hyperparameter tuning algorithms are available that are computationally expensive as the result of evaluating which samples to try next. These algorithms often have hyperparameters of their own. Bayesian optimization, random forest smart tuning, and derivative-free optimization are three examples of such algorithms.

Multivariate Testing

Multivariate testing is an extremely useful method of determining which model is best for the particular problem at hand. Multivariate testing is known as statistical hypothesis testing and determines the difference between a null hypothesis and alternative hypothesis. The null hypothesis is defined as the new model not affecting the average value of the performance metric; whereas the alternate hypothesis is that the new model does change the average value of the performance metric.

Multivariate testing compares similar models to understand which is performing best or compares a new model against an older, legacy model. The respective performance metrics are compared, and a decision is made on which model to proceed with.

The process of testing is as follows:

1. Split the population into randomized control and experimentation groups.
2. Record the behavior of the populations on the proposed hypotheses.
3. Compute the performance metrics and associated p -values.
4. Decide on which model to proceed with.

Although the process seems relatively simple, there are a few key aspects for consideration.

Which Metric Should I Use for Evaluation?

Choosing the appropriate metric to evaluate your model depends on the use case. Consider the impact of false positives, false negatives, and the consequences of such predictions. Furthermore, if a model is attempting

to predict an event that only happens 0.001% of the time, an accuracy of 99.999% can be reported but not confirmed. Build the model to cater to the appropriate metrics.

One approach is to repeat the experiment, thus performing repeat evaluations. Although not a fail-safe, this reduces the change of illusionary results. If there is indeed change between the null and alternate hypothesis, the difference will be confirmed.

Correlation Does Not Equal Causation

The phrase correlation does not equal causation is used to stress that a correlation between two variables does not suggest that one causes the other. Correlation refers to the size and direction of a relationship between two or more variables.

Causation, also known as cause and effect, emphasizes that the occurrence of one event is related to the presence of another event. It may be tempting to assume that one variable causes the other; however, in models with several features, there may be hidden factors that cause both variables to move in tandem.

For instance, smoking tobacco is a cause that increases the risk of developing a variety of cancers. However, it may be correlated with alcoholism, but it does not cause alcoholism.

What Amount of Change Counts as Real Change?

Defining the amount of change required before the null hypothesis is rejected once again depends on the use case. Specify a value at the beginning of the project that would be satisfactory and adhere to it.

Types of Tests, Statistical Power, and Effect Size

There are two main types of tests—one-tailed and two-tailed tests. One-tailed tests evaluate whether the new model is better than the original. However, it does not specify whether the model is worse than the baseline. One-tailed tests are thus inherently biased. With two-tailed tests, the model is tested for the possibility of change in two directions—positive and negative.

Statistical power refers to the probability that the difference detected during the testing reflects a real-world difference.

Effect size determines the difference between two groups through evaluating the standardized mean difference between two sets. Effect size is calculated as the following:

$$\text{Effect size} = ((\text{mean of experiment group}) - (\text{mean of control group})) / \text{standard deviation}$$

Checking the Distribution of Your Metric

Many multivariate tests use the *t*-test to analyze the statistical difference between means. The *t* value evaluates the size of the difference relative to the variation in your sample data. However, the *t*-test makes assumptions that are not necessarily satisfied by all metrics. For instance, the *t*-test assumes both sets have a normal, or Gaussian, distribution.

If the distribution does not appear to be Gaussian, select a nonparametric test that does not make assumptions about a Gaussian distribution, such as the Wilcoxon–Mann–Whitney test.

Determining the Appropriate *p* Value

Statistically speaking, the *p* value is a calculation used in hypothesis testing that represents the strength of the evidence. The *p* value measures the statistical significance, or probability, that a difference would arise

by chance given there was no real difference between two populations. It provides the evidence against the null hypothesis and is a useful metric for stakeholders to draw conclusions from.

A p value lies between 0 and 1, and is interpreted as follows:[66]

- a p value of ≤ 0.05 indicates strong evidence against the null hypothesis, thus rejecting the null hypothesis
- a p value of > 0.05 indicates weak evidence against the null hypothesis, hence maintaining the null hypothesis
- a p value near 0.05 is considered marginal and could swing either way

The smaller the p value, the smaller the probability that the results are down to chance.

How Many Observations Are Required?

The quantity of observations required is determined by the statistical power demanded by the project. Ideally, this should be determined at the beginning of the project.

How Long to Run a Multivariate Test?

The duration of time required for your multivariate testing is ideally the amount of time required to capture enough observations to meet the defined statistical power. It is often useful to run tests over time to capture a representative, variable sample.

When determining the duration of your testing phase, consider the novelty effect, which describes how user reactions in the short term are not representative of the long-term reactions. For instance, whenever Facebook updates their news feed layout or design, there is an uproar.

However, this soon subsides once the novelty effect has worn off. Therefore, it is useful to run your experiment for long enough to overcome this bias. Running multivariate tests for long periods of time are typically not a problem in model optimization.

Data Variance

The control and experimentation sets could be biased as the result of not being split at random. This may result in biases in the sample data. If this is the case, other tests can be used, such as Welch's t-test, which does not assume equal variance.

Spotting Distribution Drift

It is key to measure ongoing performance of your machine learning model once deployed. Data drifts and system development require the model to be confirmed against the baseline. Typically, this involves monitoring the offline performance, or validation metric, against data from the live, deployed model. If there is a sizeable change in the validation metric, this highlights the need to revise the model through training on new data. This can be done manually or automated to ensure consistent reporting and confidence in the model.

Keep a Note of Model Changes

Keep a log of all changes to your machine learning model with notes on changes. Not only does this serve as a change log for stakeholders, it provides a physical record of how the system has changed over time. The use of versioning software within a development environment (test/staging to live deployment) will enable software changes to automatically be noted. Versioning software provides a form of technical governance and can be used to deploy software with extensive rollback and backup facilities.

CHAPTER 6

Ethics of Intelligence

“People worry that computers will get too smart and take over the world, but the real problem is that they’re too stupid and they’ve already taken over the world.”

—Pedro Domingos

From supermarket checkouts to airport check-ins and digital healthcare to Internet banking, the use of data and AI for decision-making is ubiquitous. There has been an astronomic growth in data availability over the last two decades, fueled by, first, connectivity, and now the Internet of Things. Traditional data science teams focus on the use of data for the creation, implementation, validation, and evaluation of machine learning models that can be for predictive analytics.

In the past decade, there has been a wealth of data-driven AI and technological advance in healthcare, including the following:

- Noninvasive blood glucose levels: There have been many advances in noninvasive blood glucose management. Google developed a contact lens that determines the user’s blood glucose levels in 2016. French company PKvitality developed a small wristwatch that determines the level of glucose in the blood, and there have been advances in determining blood glucose levels through using far-infrared signals.[67]

- **Artificial pancreas:** As a combination of systems, the artificial pancreas uses two devices. The first device measures blood glucose using a sensor and communicates with a secondary device (or patch), which administers insulin to the patient. The insulin delivery system can adjust the insulin dosage according to the blood glucose levels.[68]
- **Bioprinting of skin constructs:** 3-D printing has been used to reproduce blood vessels and skin cells to facilitate wound healing for burn patients.[69]
- **Digital peer support:** Peer support communities such as Diabetes.co.uk enable peer-to-peer support have demonstrated in research studies the ability to improve qualitative and quantitative health outcomes.[70]
- **Foot ulcer detection:** Concerns such as foot ulceration, bruising, and wider diabetic foot concerns can be increasingly detected through machine learning to expedite ulceration detection, prevent amputations, determine effective treatments, and improve healing times.
- **Open source data sharing:[71]** Data repositories and public APIs have enabled data sharing between systems. Historically conservative firms are now embracing open source analytical, AI, and data management software. Within many organizations, employees are actively discouraged, or given autonomy on whether to use proprietary tools. Cost and performance are drivers toward open source data, mainly as open data sources have become more robust

and accepted. The adoption is also driven by the latest generation of data scientists and university graduates aware of the safety and capabilities of open source data.

Novel solutions to problems developed through machine learning are themselves leading to questions of morality and ethics. Currently, governance is moving at the pace of the industry itself. There are many scenarios within AI for which there are no precedents, regulations, or laws. As a result, it is paramount to consider the ethical and moral implications of creating intelligent systems.

The World Health Organization (WHO) reports that the prevalence of chronic disease will rise to 57% of the global population by 2050.[72] Unfortunately, the WHO also reports that there is a global, growing shortage of healthcare workers, which by 2035 will rise to 12.9 million.[73] Lack of professionals to provide healthcare services paints a stark image of the future with grave consequences for humanity. Healthcare professional shortages are being offset through AI, digital interventions, IoT, and other digital technologies that cannot only replace manual and cognitive working tasks but can also improve the reach, precision, and availability of healthcare. At the same time, advancements in the detection and diagnosis of diseases, genomics, pharmacology, stem cell and organ therapy, digital healthcare, and robotic surgery are expected to minimize the cost of treating illness and disease. As AI penetrates humanity's day-to-day activities, philosophical, moral, ethical, and legal questions are raised. This is amplified in healthcare, where clinical decisions can mean the difference between life and death. Even if AI can aid diagnosis of conditions or predict future mortality risk, will humans ever prefer an AI's advice over their doctor? As humankind becomes accustomed to living side-by-side to intelligent systems, there are a host of hurdles to overcome.

What Is Ethics?

Ethics or moral philosophy refers to the moral codes of conduct (or set of moral principles) that shape the decisions people make and their conduct. Morality refers to the principles that distinguish between good/right or bad/wrong behavior. Ethics in the workplace, for example, is often conveyed through professional codes of conduct for which employees must abide.

What Is Data Science Ethics?

Data science ethics is a branch of ethics that is concerned with privacy, decision-making, and data sharing.

Data science ethics comprises three main strands:

- Ethics of data
This area of data science ethics focuses on the generation, collection, use, ownership, security, and transfer of data.
- Ethics of intelligence
This area of data science ethics covers the output or outcomes from predictive analytics that data is used to develop.
- Ethics of practices[74]
The ethics of practices was proposed by Floridi and Taddeo, referring to the morality of innovation and systems to guide emerging concerns.

Data Ethics

More smartphones exist in the world than people—and phones, tablets, and digital devices alongside apps, wearables, and sensors are creating millions of data points a day. There are over 7.2 billion phones in use,

112 million wearable devices sold annually, and over 100,000 healthcare apps available to download on your mobile phone.[75] IBM reports more than 2.5 quintillion bytes (2.5×10^{18}) of data are created daily.[76] Data is everywhere. Moreover, it's valuable.

The topic of data ethics has been thrust into the public spotlight through high profile fiascos such as the Facebook Cambridge Analytica scandal. Facebook, one of the world's largest and most trusted data collection organizations, had user data harvested through a quiz hosted on its platform. Behavioral and demographic data of the 1.5 million completers of the quiz was sold to Cambridge Analytica. The data is largely considered to have been used to target and influence the outcome of the United States 2017 elections.[77] What's more concerning is that this breach of security was reported over 2 years after the initial data leak.

We are in a time where fake news can travel quicker than the truth. Society is at a critical point in its evolution, where the use, acceptance, and reliance on data must be addressed collectively to develop conversations and guiding principles on how to handle data ethically.

The ethical and moral implications of data use are vast, and best demonstrated through an example. During this chapter, we will refer to the following hypothetical scenario:

SCENARIO A

John, type 2 diabetic, aged 30, has a severe hypoglycemic episode and is rushed to the hospital.

John has been taken to the hospital unconscious for treatment. John, a truck driver, was prescribed insulin to treat his type 2 diabetes by his doctor's advice, which is most likely the cause of his hypoglycemia. Lots of data is generated in the process—both in the hospital, by healthcare professionals, and also on John's Apple Watch—his heart rate, heart rate variability, activity details, and blood oxygen saturation to look out for signs of a diabetic coma.

John's blood sample is also taken, and his genome identified. Let's suppose all this data is used, and it's useful.

John's Apple Watch was used to monitor his heart on the way to the emergency room, through which it was suspected that he has an irregular heartbeat, later confirmed with the hospital's medical equipment.

Upon waking from his hypoglycemic episode, John is pleased to learn that genetic testing does not always give bad news, as his risks of contracting prostate cancer is reduced because he carries low-risk variants of the several genes known in 2018 to contribute to these illnesses. However, John is told of his increased risks of developing Alzheimer's disease, colon cancer, and stroke from his genetic analysis.

Informed Consent

Informed consent refers to the user (or patient) being aware of what their data will be used for. Informed consent refers to an individual being legally able to give consent. Typically this requires an individual to be over 18, of sound mind, and able to exercise choice. Consent should ideally be voluntary.

Scenario A demonstrates how useful data can be given a particular context (or use case) and demonstrates the many intricacies of informed consent.

Freedom of Choice

Freedom of choice refers to the autonomy to decide whether your data is shared and with whom. This refers to the active decision to share your data with any third party. For example, should John, who has type 2 diabetes, now be required to demonstrate that his blood glucose levels are under

control and within the recommended range before being allowed to drive his truck again? A person's choice as to whether to share their data could result in a future where people are exempt from opportunities until otherwise demonstrated by their data. In an ideal world, each person should have a choice as to whether to share their data. In practice, this is neither realistic nor plausible.

There are ethical implications for John not consenting, or wanting to consent, to the healthcare team using his Apple Watch data—namely, that John's Apple Watch data belongs to John. Should the emergency response team have used John's heart rate only to ensure it was beating or was it ethical to diagnose John's atrial fibrillation subsequently? With informed consent, John would have a choice, which is a fundamental pillar of data ethics.

This is a choice that patients previously did not need to make. Before the datafication of modern life, there were fewer devices to capture such data and far less sophisticated means of predicting future events. The advances of AI and machine learning in healthcare mean that today there are two groups of people: those that seek out health information to help them plan and manage future scenarios, and those that are happy to live in a state of not knowing. As the datafication of everything continues, those that are happy to live in a state of ignorant bliss are finding less opportunity to do so.

Should a Person's Data Consent Ever Be Overturned?

In an ideal world, an individual's decision to share their data should be respected. However, as an absolute concept, this is neither realistic nor plausible. For instance, let's assume John was to decline consent to use his data. On the assumption that the emergency response team's duty is to safeguard the health of its patients, John's data helped to monitor his vital

signs, which contributed to his survival. Arguably the emergency response team would be acting unethically if they were not to use the full variety of data available to them at that given moment. Perhaps even John himself would overrule his consent if it meant increasing his chances of survival.

Precedent demonstrates freedom of choice and consent is ultimately a utopian concept. In October 2016, a man accused of the rape and murder of a 19-year-old medical student in Germany had the health data from his smartphone used against him at trial.[78] The suspect, who was identified by a hair discovered at the crime scene, refused to give police the PIN code to his smartphone. Police enlisted the help of a cyberforensics firm in Munich who broke into the device. Data on the suspect's iPhone recorded his steps and elevation, which police analyzed. Police suggested the suspect's elevation (stair climb) data could correlate to him dragging his victim down a riverbank and climbing back up. As well as locating the suspect's movements, the phone also suggested periods of strenuous activity, which included two peaks the onboard smartphone app put down to climbing stairs. Police investigators mimicked how they believed the suspect disposed of the body and demonstrated the same two peaks of stair climbing detected on the suspect's iPhone.

Public Understanding

The Facebook Cambridge Analytica fiasco went on to highlight just how unaware the public are on the topic of data privacy. The US Senate's interrogation of Facebook CEO Mark Zuckerberg illustrated just how ignorant and unaware the public was on the topic of technology. Senators were puzzled as to how Facebook made money as a free platform, referred to sending e-mails through WhatsApp, and queried as to whether encrypted messages could ever be used to provide targeted advertising.[79] The misunderstandings and ignorance demonstrated by critical stakeholders on the topic of data governance were astounding.

Increased public awareness and understanding are required to educate the public on the use of data and to empower people to decide how, where, and by whom their data can be used.

Who Owns the Data?

People generate thousands of data points daily. Nearly every transaction and behavior creates a data trail left through devices such as smartphones, televisions, smartwatches, mobile apps, health devices, contactless cards, cars, and even fridges. However, who owns the data? The topic of data ownership is an example of a complex first-time problem that has facilitated the development of international policy and governance. Historically, user data is owned by companies rather than the individual.

EHRs enable machine learning-based predictive analytics that will eventually enable providers to provide enhanced levels of care. Although it seems unlikely a patient would not want to not share their health data for improved morbidity and mortality, one would reasonably assume it would become more difficult to withdraw such privileges, much akin to closed-circuit television (CCTV) where consent is not often noted or obvious, and is widely assumed to have utilitarian benefits over the long term. Many platforms prevent users from accessing their full range of services if data sharing is disabled. This is driven by the organizational need to store data in a robust and central repository for safety, governance, and improvement.

Take the example of a patient who uses a connected blood glucose device and mobile app to track and record their blood glucose. Blood glucose data is delivered from the user's blood glucose meter to her mobile application. Although the data would be presented on the user's phone, the data is held in the mobile application provider's database, subject to the terms and conditions of use. Many of today's web sites, mobile apps, connected devices, and health services state that data can and will be used in an anonymized and aggregated format, or in some instances in an identifiable format, by the organization providing the service and also typically with selected partners.

Data collection, usage, and sharing have become a fundamental part of supporting the data-driven quality improvement processes.

Types of shared data include the following:

- Anonymized data

Anonymized data is data that has identifiable features removed. Identifiable features are features that enable someone to recognize the person that the data has come from. For example, an oncology ward's spreadsheet of patients would be anonymized if the name, date of birth, and patient number were removed.

- Identifiable data

Identifiable data refers to data that can be used to identify an individual. For instance, an oncology ward's spreadsheet of patients would be identifiable if it reported the name, date of birth, and patient number.

- Aggregate data

Aggregate data refers to data that has been combined and where a total is reported. Following the oncology example, if the ward's spreadsheet covered ten patients, aggregate data reporting could include the male-to-female split or age brackets of patients. Data is cumulatively reported for the population within the dataset.

- Individualized data

Individualized data is the opposite of aggregated data. Instead of data being combined, data is reported for each person within the dataset. Individualized data does not have to be identifiable.

Concerns over data privacy have spurred a global response. In May 2018, GDPR (or the General Data Protection Regulation) came into force across Europe. GDPR legislation governs organizations in how they use and share user data.[80] Because of GDPR, organizations have been forced to collect the opt-in consent of their memberships and declare to users how they would use user data and with whom they should share it. The GDPR legislation places data control firmly in the user's hands. Users of data-driven systems are now able to exercise their rights in seeing the data held on them, with whom their data is shared and why, the right to be forgotten, and the right to be deleted. GDPR also defined labels for those collecting and processing data:

- Data controller: The data controller refers to the person or organization that controls, stores, and makes use of the data.
- Data processor: The data processor refers to the person or organization who processes data on behalf of a data controller. Based on this definition, agents such as calculators could be considered data processors.

For those working with data, it is useful to understand the distinction between data processor and controller, and the responsibilities of each party.

GDPR demonstrates how vital data safety has become. GDPR is tightening data access, security, and management in an ambitious attempt to protect the EU's 500 million citizens. It has single-handedly repositioned control with users. The maximum fine for the worst offenders against GDPR regulations is set at either €20 million (£17.6m) or 4% of global revenues for larger organizations.[81] Unfortunately, regulations affect all businesses and organizations, and a flurry of reconsent requests flooded the Internet. There is much confusion as to how GDPR laws are applied, as demonstrated when Mark Zuckerberg was summoned to the European

Parliament in 2018. Just like in the United States, key figures were ignorant and unaware as to how data and connectivity work in the twenty-first century.

The topic of a user's right to be forgotten within machine learning has posited several compelling ethical questions. For instance, if John's data were to be used in a machine learning algorithm to predict the likelihood of severe hypoglycemia, and if John were to request his data to be deleted, it would be close to impossible to decouple John's data from the machine learning model that is learned. Should John's consent be overridden if there is a utilitarian benefit?

Further still, if John's data were to be for some reason valuable in developing an algorithm to diagnose and predict disease, it could be claimed that it is unethical for John to refrain from sharing his data or requesting his data to be removed from such an algorithm. If John were to be the only patient in the world to have a genetic defect, or if his data was for some reason useful to progress medical understanding, it is evident as to how and why John's data would be useful to humankind.

What Can the Data Be Used For?

Data is already used to fuel a variety of decisions. Employers have used psychometric and parametric testing for decades to understand potential candidates. Nowadays, employers also look at supporting sources of data. Employers often scour the social media profiles of prospective employees for concerns of reputational risk. In the same vain, history is plagued with people who have lost their jobs as the result of (deemed inappropriate) social media use.

What an individual's data can be used for raises ethical and moral apprehensions. The concepts of identity and free will are challenged by the notion that an individual's data could be used to make decisions with direct ramifications that may be unbeknownst to that person.

Car insurance is an industry that has evolved to optimize the use of data. Historically, car insurance premiums have been grounded in accident claim data, which is generalized for segments of the population. The advent of the black box enables more precise premiums based on a variety of demographic and behavioral factors—such as the age of the driver, times that the car is used, the speed of driving, and frequency of erratic acceleration among others. This data can also be used to reduce the cost of car insurance premiums should the driver be sensible and enables the increase in car insurance premiums should the driver not follow the stipulations set out by the insurer. Others see this constant real-time data analysis (and subsequent feedback) as “big brother.”

Similarly, life and health insurance are beginning to follow in the footsteps of car insurance. Life and health insurers are increasingly providing incentivized products that reward positive behaviors and chastise negative behaviors. Life insurer Vitality, for instance, offers a product whereby it provides an Apple Watch to its members to encourage positive and healthy behaviors. Similarly, British organization Diabetes Digital Media provide evidence-based digital health interventions to insurers and bill payers worldwide to incentivize wellness.[82]

Individuals should be aware that data could also be used for cynical, more sinister purposes. For instance, patient data could be used to decline an individual from a treatment, operation, or opportunity. The public has been incredibly trusting of data use, and robust data governance is required to prevent its future misuse. GDPR is regarded as the first step toward improved understanding and data governance. GDPR is considered by some to be as revolutionary to data science as the advent of the Internet itself.

Privacy: Who Can See My Data?

Conversations of data ownership naturally lead to who can see your data. The key is in ensuring only approved services and organizations have access to your data. For instance, it is unlikely you would want a life insurance company to be given your medical data without your consent, particularly if there was data that could influence your coverage.

Data sharing between applications is commonplace, with APIs enabling accessibility and faster connectivity between independent services. For instance, users can import nutrition data from apps like MyFitnessPal into their diabetes management or fitness apps. These services often replicate user data among a variety of independent architectures, which leads to concerns over managing approved data access. Applications that enable data integration must also provide the facility to decouple patient data. Systems must be able to verify imported data for data governance, auditing, and patient safety.

The debacle surrounding Facebook and Cambridge Analytica went on to demonstrate that even if one trusts an aggregator of data, it is possible that third parties are engaging with that very data without your knowledge. Facebook requested Cambridge Analytica to delete its Facebook user data. Even though Cambridge Analytica agreed to the request and said they had deleted the user data, they were later demonstrated not to have deleted the data.[83] This act of deception leads to questions of accountability, the ramifications of data leaks, and who is accountable in situations of third-party data leaks.

Moreover, even if data is anonymized, it may not guarantee privacy. Netflix recently published 10 million movie rankings from 500,000 members as part of a challenge to improve the Netflix recommendation system. Although data was anonymized to remove personal details of members, researchers at the University of Texas were able to de-anonymize some of the Netflix data through comparing similar rankings and time

stamps with public information in IMDb (Internet Movie Database). There are natural security problems with anonymous data—mainly that it does not mean that you are anonymous.

When it comes to sharing data, it is useful to distinguish between people and patients. The general public's attitude toward data use in the realm of healthcare is far more welcoming than its attitude toward nonmedical data use. A survey of 5,000 people by Diabetes.co.uk reported that 56% of people are hesitant to share their data without good reason. When a group of patients with type 2 diabetes was asked a similar question about whether their data could be collected and used for further research, 83% of patients opted to share their data.[84] Patients appear to understand or are at the very least hopeful that, through sharing of health data, they are progressing medical understanding and treatment. People, however, are far more skeptical of beneficial data use. Confidentiality is a core tenet of data ethics.

How Will Data Affect the Future?

The sharing of patient data and aggregation of big datasets is being used to enhance diagnosis, treatment, and care. As the types and quality of data improve, healthcare's precision will become more precise in the areas discussed next.

Prioritizing Treatments

Big data medical sets can enable predictive analytics to determine optimal treatment pathways for various segments of the population. There is the potential to prioritize only those who are likely to see efficacy in the treatment, which begs the question as to how best to support those that do not.

Determining New Treatments and Management Pathways

The analysis of real-world experience data, clinical studies, randomized clinical trial (RCTs), and pharmacological data is facilitating treatment and management discovery. Digital health interventions have been demonstrated to reverse type 2 diabetes, reduce the number of epilepsy seizures. Data has the potential to develop new pharmacological interventions and disrupt traditional treatment paradigms.

More real-world evidence

The use of real-world evidence (RWE) at scale from patient communities, digital education programs, and health tracking apps is being increasingly demonstrated to improve the self-management landscape. A concern with real-world evidence has typically been lack of academic robustness. However, RWE is being increasingly used in research to determine population usage and benefits. Real-world evidence is a vital part of AI's ethical journey, and there must be trust between patient and provider.

Enhancements in Pharmacology

The precision audiences required for RCTs and academic studies can be identified quicker and easier through digital platforms. This has benefits, such as quicker project recruitment times, greater potential, and multisource comparison. Real-world data is being used to develop better and more effective drugs.

Optimizing Pathways Through Connectivity—Is There a Limit?

Machine learning's primary use in the short term involves data analysis and predictive analytics.

Security

There are significant privacy concerns that come with a truly unified system. What if such a system were to be compromised? Do people, or patients, want a truly unified system, or is the possibility to use it for malintent greater than the potential good? The consequences of vulnerabilities—whether security, operational, or technical—is amplified in unified systems.

For years, large companies have talked about ways to combine and analyze the giant repositories of data they separately gather about people: location data from smartphones, financial data from banks, relationship data from social networking apps, and search data from browsers—to build a complete picture of a person's behavior.

Facebook reportedly offered to match the data that hospitals had about individual patients with social information gleaned from the social networking site, such as how many friends the person has, or whether they seem to engage with others on the site. The company paused the project after privacy concerns were raised by the Cambridge Analytica scandal.

Further still, there are security concerns around IoT devices, which now range from thermometers, cars, and washing machines to blood glucose meters, continuous glucose monitors, and insulin pumps. The requirement of network connectivity to operate as a smart device leaves all devices prone to vulnerability. Distributed denial of service attacks (DDoS) is used on IoT devices to break in and leave malicious code to develop botnets, leak data, or otherwise compromise the device. WeLiveSecurity

reported on 73,000 security cameras with default passwords in 2016.[85] A basic but key to learning from this is that it is always advisable to change credentials that use default passwords. DDoS attacks are nothing new, but the extent of vulnerability is only being uncovered through research and first-time problems, or security breaches. Research students from the University of Florida were able to compromise Google's Nest thermostat in under 15 seconds.[86]

DDoS attacks can be crippling for any organization. Develop mitigation procedures and ensure network infrastructures allow visibility of traffic entering and exiting the network. It is good practice to develop a DDoS defense plan, which should be kept and regularly updated and rehearsed.

Ethics of Artificial Intelligence and Machine Learning

A primary application of machine learning in healthcare involves patient diagnosis and treatment. AI models are deployed to help physicians diagnose patients, especially in cases involving relatively rare diseases or when outcomes are hard to predict.

Imagine a future where doctors know exactly how many times you've eaten fast food in the last month, and that it is connected to your medical record. That data is used to suggest what foods to eat, comparing your health record to hundreds or thousands of other people to identify who is like you. Moreover, imagine this could directly affect life insurance or health insurance, with possible rewards for positive behavior and avoiding bad foods. Imagine a system that was able to predict your likelihood of developing any particular disease in real time.

The ethics of machine learning refers specifically to the questions of morality surrounding the outputs of machine learning models that use data, which come with their ethical concerns.

Machine learning has already been used to develop intelligent systems that have been able to predict mortality risk and length of life from health biomarkers. AI has been used to analyze data from EHR to predict the risk of heart failures with a high degree of certainty.

Moreover, machine learning can be used to determine the most effective medication dosage learning on patient real-world and clinical data, reducing healthcare costs for the patients and providers. AI can not only be used in determining dosage but also in determining the best medication for the patient. As the genetic data becomes available, medications for conditions such as HIV and diabetes will accommodate for variations among races, ethnicities, and individual responses to particular drugs. Within the same data, medication interactions and side effects can be tracked. Where clinical trials and requirement for FDA approval look at a controlled environment, big real-world data provide us with real-time data such as medication interactions and the influence of demographics, medications, genetics, and other factors on outcomes in real time.

As the limitations of technology are tested, there are ethical and legal issues to overcome.

Machine Bias

Machine bias refers to the way machine learning models exhibit bias. Machine bias can be the result of several causes, such as the creator's bias on data used to train the model. Biases are often spotted as first-time problems with subtle and obvious ramifications. All machine learning algorithms rely on a statistical bias to make predictions about unseen data. However, machine bias reflects a prejudice from the developers of the data.

The capabilities of AI regarding speed and capacity of processing far exceed that of humans. Therefore, it cannot always be trusted to be fair and neutral. Google and its parent company, Alphabet, are leaders when it comes to AI, as seen with Google's Photos service, where AI is used to identify people, objects, and scenes. But it can still go wrong, such as when

the search engine showed insensitive results for comparative searches for white and black teenagers. Software used to predict future criminals have been demonstrated to show bias toward black people.[87]

Artificially intelligent systems are created by humans, who are biased and judgmental. If used correctly, and by those who positively want to affect humanity's progress, AI will catalyze positive change.

Data Bias

Bias refers to a deviation from the expected outcome. Biased data can lead to bad decisions. Bias is everywhere, including the data itself. Minimize the impact of biased data by preparing for it. Understand the various types of bias that can creep into your data and effect analysis and decisions. Develop a formal and documented procedure for best practice data governance.

Human Bias

For as long as humans are involved in decisions, a bias will always exist. Microsoft launched an AI chatbot named Tay in 2017. Tay interacted with other Twitter users and learned from its interactions with others. Once the Twittersphere seized hold of Tay, trolls steered the conversation from positive interactions to interactions with trolls and comedians. Within hours, Tay was tweeting sexist, racist, and suggestive posts. Sadly, Tay only lived for 24 hours. This experiment raises the question of whether AI can ever really be safe if it learns from human behavior.[88]

Intelligence Bias

Machine learning models are only as good as the data they are trained on, which often results in some form of bias. Human thinking AI systems have demonstrated bias amplification and caused many data scientists to discuss the ethical use of AI technology. Early thinking systems built on

population data showed significant signs of bias regarding sex, race, social standing, and other issues.

An infamous example of algorithmic bias can be found in criminal justice systems. The Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) algorithm was dissected in a court case about its use in Wisconsin. Disproportionate data on crimes committed by African Americans were fed into a crime prediction model, which then subsequently output bias toward people from the black community. There are many examples and definitions of biased algorithms.[89] Algorithms that assess home insurance risk, for instance, are biased against people who live in particular areas based on claim data. Data normalization is key. If data is not normalized for such sensitivities and systems not properly validated, humanity runs the risk of underrepresenting and skewing machine learning models for minorities and underrepresenting many groups of people.

Removing bias does not mean that the model will not be biased. Even if an absolute unbiased model were to be created, we have no guarantee that the AI won't learn the same bias that we did.

Bias Correction

Bias correction begins with the acknowledgment that bias exists. Researchers began discussions on machine learning ethics in 1985 when James Moor defined implicit and explicit ethical agents.[90] Implicit agents are ethical because of their inherent programming or purpose. Explicit agents are machines given principles or examples to learn from to make ethical decisions in uncertain or unknown circumstances.

Overcoming bias can involve post-processing regarding calibration of our model. Classifiers should be calibrated to have the same performance for all subgroups of sensitive features. Data resampling can help smoothen a skewed sample. But, for many reasons, collecting more data is not very easy and can cause budgetary or time problems.

The data science community must actively work to eliminate bias. Engineering must honestly question preconceptions toward processes, intelligent systems, or how bias may expose itself in data or predictions. This can be a challenging issue to tackle, and many organizations employ external bodies to challenge their practices.

Diversity in the workplace is also preventing bias from creeping into intelligence. If the researchers and developers creating our AI systems are themselves lacking diversity, then the problems that AI systems solve and training data used both become biased based on what these data scientists feed into AI training data. Diversity ensures a spectrum of thinking, ethics, and mind-sets. This promotes machine learning models that are less biased and more diverse.

Although algorithms may be written to best avoid biases, doing so is extraordinarily challenging. For instance, even the motives of the people programming AI systems may not match up with those of physicians and other caregivers, which could invite bias.

Is Bias a Bad Thing?

Bias raises a philosophical question on the premise that that machine learning assumes that biases are generally bad. Imagine a system that is interpreted by its evaluators to be biased, and so the model is retrained with new data. If the model were to output similarly biased results, the evaluation might wish to consider that this is an accurate reflection of the output and hence require a reconsideration of what bias is present. This is the beginning of a societal and philosophical conflict between two species.

Prediction Ethics

As advanced machine learning algorithms and models are developed, more accurate and reliable conclusions will be achieved in a short space of time. Technologies are being used currently to interpret a variety of

images, including those from ultrasound, magnetic resonance imaging (MRI), X-rays, and retina scans. Machine learning algorithms can already effectively identify potential regions of concern on images of the eye and develop possible hypotheses.

It is key to ensure trust in your goals, data, and organization through good governance and transparency. This is fundamental for AI going forward.

Explaining Predictions

As AI algorithms become smarter, they also become more complex. Remaining ignorant about the construction of machine learning systems or allowing them to be constructed as black boxes could lead to ethically problematic outcomes. If an agent were discovered to be predicting incorrectly, it would be a daunting task discovering the behavior that caused an event that is hidden away and virtually undiscoverable.

Interpretability of both data and machine learning models is a critical aspect of intelligent systems. This not only ensures the model integrity but also that it is attempting to solve the correct problem. Users of solutions that embed data science will always prefer experiences where they are understandable and explainable. Data scientists can also use interpretability metrics as the basis for validation and improvement.

Machine learning black boxes could harbor bias, unfairness, and discrimination through programmer and data choices to never be known. Neural networks are an example of a typically unexplainable algorithm. The backpropagation algorithm's computed values cannot be explained. As AI develops in its accuracy in resembling humans, there will be a greater requirement to ensure AI isn't picking up the bad habits of humans.

Protecting Against Mistakes

Intelligence comes from learning, whether you are a human or machine. And intelligent machines, just like humans, learn from mistakes. Data scientists will typically develop machine learning models with a training, testing, and validation phase to ensure systems are detecting the correct patterns within a defined tolerance. The validation phase of machine learning model development is unable to cover all possible permutations of parameters that may be received in the real world. These systems can be fooled in ways that humans wouldn't be. Governance and regular auditing is required to ensure that AI systems perform as intended and that people cannot influence the model to use it for their own ends.

The incorrect classification of predictions can lead to scenarios where the outcome is a false positive or false negative. The impact of both should be considered in the context of your domain. Take for instance an incorrect diagnosis of breast cancer. In a false positive scenario, a patient would be informed they had breast cancer when they did not. Being informed that this classification was incorrect will come to some relief to the patient. A false negative would result in a patient's disease progression and the eventual correct re-diagnosis. The mental and physical trauma as a result of a false negative prediction should be considered. Patients should always be informed of the level of accuracy of results.

Information governance for systems that will be used by patients (patient-facing), whether predictive or otherwise, should provide robust risk mitigation procedures for mistakes in outcomes. Emotional or psychological patient support should be considered for those experiencing trauma as the result of the incorrect diagnosis.

As well as acting as a catalyst for innovation, the speed at which agile digital technologies are rushed to the market is also AI's biggest pitfall. Technology company LG infamously unveiled an AI bot at CES 2018 (Consumer Electronics Show, an annual trade show organized by the Consumer Technology Association) that ignored the presenter's

instructions. The AI bot, which was marketed as providing innovative convenience, failed to respond to any of its master's comments, was either malfunctioning, or chose to ignore the commands.[91] Whether in healthcare or not, AI system performance metrics should be transparent and audited regularly. The cost, particularly in healthcare, is too high not to. AI systems in healthcare have been demonstrated to fail exorbitantly. In the UK, an NHS breast cancer screening system failed to appropriately invite women for screening, with observers claiming up to 270 females may have died as a result.[92] In this particular case, the organization responsible for running the system placed blame with a contractor. The ethical implications and public perception of such mistakes are enormous.

Validity

There is a requirement to ensure the validity of machine learning models over time to ensure the model can be generalized and generalizations are valid. Regular testing and validation of models are essential to maintaining integrity and precision of your machine learning model. A suboptimal predictive analytics model will provide unreliable results and damage integrity.

Preventing Algorithms from Becoming Immoral

Algorithms can, and do, already act immorally. To date, AI algorithms have mainly performed in unethical ways due to design. This has been demonstrated very well outside healthcare by organizations such as Uber and Volkswagen. Uber's Greyball algorithm attempted to predict which passengers were undercover police officers and used it to identify and deny transport.[93] Volkswagen's algorithm allowed vehicles to pass emission tests by reducing nitrogen oxide emissions during the test phase.[94] Both organizations were internationally condemned for their public

deception and lack of transparency. Precedents such as these from the world's largest companies demonstrate that internal and external auditing is required to validate the integrity and ethics of algorithms and indeed their organizations.

There is a genuine concern that AI may learn to act immorally not only from its creators but also from its experience.

The Ecole Polytechnique Fédérale of Lausanne's Laboratory of Intelligent Systems in Switzerland conducted a project that monitored robots designed to collaboratively search for positive resources and ignore dangerous items.[95] Robots were designed as genetic agents equipped with sensors and a light that was used to flag the identification of a positive resource, which was finite in number. Each agent's genome dictated its response to stimulus and was subject to hundreds of generations of mutations. Agent learning was reinforced by the agent receiving positive marks for identifying positive resources and negative marks for its proximity to poisonous items. The top 200 highest performing genomes were mated and mutated at random to produce the next generation of agents. In their first generation, agents switched their lights on when they discovered a positive resource. This enabled other agents to find the positive resource. Due to the limited number of positive resources, not all agents were able to benefit. Overcrowding meant that some agents would also be distanced from the positive resource they initially found. By the 500th generation, the majority of agents had evolved to keep their light switched off when they discovered a positive resource, and a third of agents evolved to act in a way that was the exact opposite of their programming. Some agents had grown to identify lying agents that had an aversion to the light. Agents that were initially designed to cooperate eventually ended up lying to each other due to scarcity.

These concerns are amplified when applied to incentivized clinical decision-support systems that could generate increased profits for their

architects or providers—a system that was recommending clinical tests, treatment, or devices in which they hold a stake, or by altering referral patterns.

In healthcare, this scenario is very concerning. All machine learning models regardless of their use must be governed and verifiable. A machine learning model that was incentivized to make decisions should also adhere to robust standards of transparency, morality, and scrutiny.

Unintended Consequences

The adoption of AI in medicine is becoming more commonplace; and as a result, first-time problems or unintended consequences will govern the media perception and direction of AI ethics. There are very few technologies that have been embedded into healthcare without unintended, or adverse, effects. The question quickly comes to how data scientists, on behalf of humanity, can mitigate these risks.

The ethics of AI and unintended consequences have been significantly directed by The Three Laws, published by science fiction writer Isaac Asimov in 1942.[96] Asimov's Laws, found in the "Handbook of Robotics, 56th Edition, 2058 A.D.," were designed as a safety feature to prevent robots from harming humans:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey orders given by a human being unless it conflicts with the First Law.
3. A robot must protect its existence as long as such protection won't conflict with the First or Second Law.

In Asimov's stories, man-like machines behave in counterintuitive ways. These acts are the unintended consequence of how the agent applies Asimov's Laws to its environment. Only 70 years later, Asimov's science fiction fantasy is eerily turning into reality. In 2012, South Korea published a Robot Ethics Charter to prevent ills and malintent;[97] and the IEEE and British Standards Institute have both published best practice guides for ethical agent engineering.[98] Best practice guides to designing ethically sound agents are typically grounded in Asimov's Laws.

Fundamentally, intelligent agents are made of binary code and do not contain or obey Asimov's Laws. Their human architects are responsible for implementing the laws and reducing the risk of unintended consequences.

AI is unlikely to become evil and turn on humanity in the way depicted by Hollywood. Rather, a lack of context could lead to AI making unintended and disastrous actions. For instance, an intelligent agent that was tasked with eradicating HIV in the global population could eventually come to some conclusion that to achieve its objective; it should kill everyone on the planet. It is easier to frame AI intelligence in a negative setting. Without careful management, an AI agent's utility function could allow for potentially harmful scenarios. There is a limited foundation to suppose that an AI agent would have such an extreme adaptation, yet it is worth consideration. There are undoubtedly many positive unintended consequences to AI that could save humanity from itself.

How Does Humanity Stay in Control of a Complex and Intelligent System?

Over millions of years, humans have combined ingenuity and intelligence to create methods and tools to control other species. As a kind, humans have evolved to dominate, driven primarily by the capacity to learn from the mistakes of themselves or others. Through this, humans have come

to develop tools to master bigger, faster, stronger animals and techniques such as mental or physical training to achieve optimal performance for such tasks.

What will happen when AI is more intelligent than humans? Human-made AIs are already able to surpass human cognition in niche areas. AlphaGo Zero, the reinforcement learning agent developed by Alphabet, acted as its teacher to master the game of Go.[99] AlphaGo was able to learn itself, without the aid of humans or historical datasets over thousands of game interactions. Are humans doomed to be mastered by an evolving AI race? This concept is known as the singularity—the point at which AI intelligence surpasses that of humans. It cannot be expected that an AI whose intelligence surpasses humans can be switched off, either. A reasonably intelligent agent may anticipate this action and potentially defend itself. With AI that is more intelligent than humans, particularly if learning itself (with its hyperintelligence), humankind is no longer able to predict outcomes and must accordingly be prepared for the situation.

Intelligence

Intelligence refers to the outcomes of AI models and how they are used. Intelligence powers the ads you see, the apps you download, the content you see on the Internet, the cab you hire, and the price you pay for items such as loans and mortgages.

The ethics of intelligence posits questions such as whether an autonomous vehicle should protect its driver or others? Should a car prioritize protecting its passengers or instead protect other drivers and the public? What if an incident were to cause loss of life—what should be the approach then? This may become less of a problem once autonomous vehicles are prevalent and fewer humans are involved. Until then, if an autonomous vehicle on autopilot and without human steering does have

an accident, who should be liable? Court cases involving autonomous vehicle accidents have demonstrated that drivers are often using autopilot on impact.

Assessments have expressed that humans do not want a fair, utilitarian approach and want their vehicle to protect the driver. Regulation on topics such as these will be instrumental in the long-term adoption of autonomous vehicles.

A significant application of machine learning is NLP. Virtual assistants such as Alexa have revolutionized methods of agent engagement. As the use of AI becomes embedded with everyday life, humans must not be assumed to be vigilant and aware, and organizations should take steps to assure the public. The vocal skills of Google's Duplex bot were demonstrated at its developer conference in 2018 where it was shown booking a hair appointment. The onstage unveiling involved the Duplex software conducting a conversation with a hair salon receptionist. The computer-generated voice used pauses, colloquialisms, and circumlocutions usually present in human speech. The voice was under control of Google's DeepMind WaveNet software that has been trained using lots of conversations, so it knows what humans sound like and can mimic them effectively. Although welcomed by some, many were concerned at the deliberate deception of a human by AI. This has amplified the public's voice that humans demand explicit confirmation when engaging with AI.[100]

Natural language can be separated into two components: style and content. Analysis of communications such as engagements, e-mails, or messages can easily identify sentiment and opinion. Negative terms, for example, could be identified to understand the state of someone's emotional health. The implications of this, taking place publicly or privately, are vast. Analysis of social media profiles and other unstructured sources of data could lead to employers or insurers making predictions about people based on engagement, communication, or sentiment data available for mining.

Health Intelligence

Health intelligence refers specifically to the intelligence developed through healthcare AI. There are a variety of areas that health intelligence is being used in industry, driven by improvements in patient health, costs, and resource allocation:

- Healthcare services—patients are increasingly engaging with predictive analytic services to diagnose disease. For instance, diabetic retinopathy is being detected through AI systems developed through a partnership between Alphabet and Great Marsden hospital. Similarly, British organization Diabetes Digital Media developed a foot ulceration detection algorithm for earlier and quicker referral to podiatry clinics.
- Pharmacology—new drugs are being discovered through the use of learning on real-world data and patient profiles. This is enabling pharma to develop new and more precise medications.
- Life and health insurance—patients with health conditions (or risks) such as type 2 diabetes and prediabetes, who would traditionally receive loaded premiums, can use digital health interventions to manage and improve their condition: and as a result of their sustained engagement, receive reduced premiums to incentivize wellness. The benefits to life and health insurers are vast. Not only are insurers able to engage with more of the population, through connecting insurance with improved wellness, but insurers also can save money through fewer and reduced claims, reduced paramedical and medication costs, improve their risk profile, and optimize risk algorithms and underwriting.

Health intelligence cannot be developed in an environment where the public's ethical AI requirements are ignored, such as with personal data. The ethics of AI should be integral to the development of systems capable of health intelligence, identifying the risks of agents before they are realized. Health intelligence ethics are integral to understanding the intention of data and use of outcomes in medicine.

Who Is Liable?

AI's application in healthcare is an exciting opportunity to improve patient care and cost savings within a short period. Accurate, timely decisions influence the diagnosis, treatment, and outcomes of patients with any number of illnesses. Healthcare professionals can analyze a finite number of images, tests, and samples; and clinical decisions are still prone to human error. With AI, clinical decisions can be made with varying confidence through the analysis of infinite samples in near real time. The medical and technological limitations of AI appear easier to overcome than the potential moral and legal issues. If a breast cancer diagnosis algorithm falsely predicts breast cancer or an algorithm fails to identify signs of diabetic retinopathy on an eye scan, who should be to blame?

There are three liability possibilities in the diagnosis of a patient's disease:

- A human doctor decides a patient diagnosis, with no assistance from an external agent. A human doctor may be very accurate when symptoms are evident and often spots less obvious concerns at first glance. In this case, the liability is always with the doctor.
- An intelligent agent predicts the patient diagnosis, with a 99% accuracy. Seldom will a patient be misdiagnosed, but errors such as death are neither the liability of the agent or human. If an AI was to make such a mistake, how would a patient raise their issue?

- A human doctor is assisted by an intelligent agent. In this case, it may prove harder to identify responsibility, as the prediction is shared. Even if the human's decision was final, it could be argued that the countless experience of the AI agent would be a key factor in influencing a decision.

Liability for an AI agent's behavior can be understood as belonging to the following:

- The organization that developed the AI agent.
- The human team that designed the AI agent is responsible for any unexpected functioning or inaccurate predictions.
- The AI agent is responsible for any unexpected behaviors itself.

AI agent development is often the result of numerous engineers and collaborators, and thus holding the developers to account is not only difficult to manage but may put off potential engineers from joining the industry. Holding the AI development organization to account sounds like the best route to ethical adherence.

How do we ensure AI systems will not overturn humans? What regulations govern our safety? Friedler and Diakopoulos suggest that five core principles are used to define organizational accountability:[101]

- **Auditability:** External bodies should be able to analyze and probe algorithm behavior.
- **Accuracy:** Ensure good clean data is used. Identification of evaluation metrics, regular tracking of accuracy, and benchmarking should be used to calculate and audit accuracy.

- **Explainability:** Agent decisions should be explainable in an accessible manner to all stakeholders.
- **Fairness:** Agents should be appraised for discrimination.
- **Responsibility:** A single point of contact should be identified to manage unintended consequences and unexpected outputs. This is similar to a Data Protection Officer, but specifically for AI ethics.

First-Time Problems

Real-world data is already demonstrating that type 2 diabetes management has been misguided for the past 50 years. Who should be liable when real-world experience demonstrates legacy treatments have been misled?

Supervised and unsupervised models developed on patient data is creating new concepts in healthcare that are defining AI ethics. First-time problems are a form of unintended consequence. Data can now be used within AI to detect risks of disease, such as type 2 diabetes, hypertension, and pancreatic cancer. If an application was able to tell you about a terminal disease you could do nothing about, would you want to know? Would you even want to know that there was a choice to know? Services like 23andMe give feedback on disease risk based on the patient's genetic profile. For instance, if someone were to be informed they had a reduced risk of developing certain diseases, there may be some individuals that have a propensity for riskier behaviors. If a patient is told their susceptibility to lung cancer is lower than average, would this lead to a higher risk of them becoming smokers or other risky behaviors? Similarly, knowing the risk of developing particular conditions could lead to mental and emotional consequences. The implications on human psychology and behavior of knowing disease risk will be discovered as time and evidence develops.

Defining Fairness

AI is changing the way we interact and engage with products and services. As accessibility to data grows, how do we ensure AI agents treat people justly? Recommendation systems, for instance, can significantly affect our experience. But how do we know they are fair? Are we being biased by the choices we are given? A machine learning model cannot be fair if there is not a clear definition of fairness. There are many definitions of fairness and collaboration between social scientists and engineers, and AI researchers are required to determine a clear recognition of fairness. Fairness dictates that data misuse should have public, legal, and ethical consequences.

How Do Machines Affect Our Behavior and Interaction

AI's rapid evolution is prompting humankind to evaluate just what it means to be a human.

Humanity

As AI bots have become better at modeling conversations and engagements with humans, so too have their prevalence. Virtual assistants such as Alexa, Siri, and Google Home are in the hands of over 100 million people—and voice-based telephony agents are commonplace. Not only are agents able to mimic conversational language, but emotion and reaction modeling is more sophisticated and believable than ever. In 2015, a bot named Eugene won the Turing Challenge, which asks humans to rate the quality of their communication with an unknown agent and to guess whether it is a human or an AI agent. Eugene was able to disguise itself as a human to more than half of its human raters.[102] The blurring between humankind and robot-kind is imploring exploratory interspecies

relationships. In 2017, Chinese AI engineer Zheng Jiajia married robot Yingying, a robot spouse he created to ease the pressure of marriage.[103] Human-replacement sex robots are now common, with connected AI robots developed with human-like skin and warmth, as well as safety measures to mitigate electrical shocks or worse.

Whether applied to banking, healthcare, transport, or anything else, humanity is at the beginning of a juncture where humans engage with AI agents as frequently as though they too are humans. Where humans have limits as to the amount of attention, kindness, compassion, and energy they can expend, AI bots are an infinite resource for developing and maintaining relationships.

Behavior and Addictions

Our behaviors are already influenced and manipulated by technology. Web landing pages, shared links, and user experiences are optimized through multivariant testing. This is a crude algorithmic approach to capture human attention. In the 2017 US Elections, Donald Trump's campaign famously used and tested more than 50,000 ad variations daily to micro-target voters.[104] Human experimentation has been taking place for centuries, so this is nothing new. However, there is now the ability to influence choice and freedom like never before. Organizations have a moral responsibility to analyze the impact of AI and safeguard particular groups of people, including vulnerable adults and children, from manipulative tools like nudging or social proofing. Codes of conduct for behavioral experimentation within technological experiments is required.

Technology addiction is becoming a growing concern. Robert Lustig, a professor at the University of Southern California, discovered that the human brain responds to technology in a similar way to other addictive substances.[105] Internet addiction has been documented in a variety of countries; and adolescents appear to be more at risk, partly due to the prefrontal cortex being the last part of the brain to develop. According

to the charity Action on Addiction, one in three people is addicted to something.[106] In 2015, a 38-year-old man was found dead after playing video games for five days in an Internet cafe in Taipei.[107]

Real-world and academic evidence demonstrates that technology addiction is a mounting issue. AI has the potential to improve productivity and discovery. However, humankind is responsible for taking measures to prevent AI leading to destructive, digital addictions that make us more insular, overdependent, and lethargic.

Economy and Employment

What happens after the end of jobs? Supermarkets already use automated checkouts, and even McDonalds have already replaced the human-to-human food ordering procedure with a digital interface.[108] The hierarchy of labor is concerned primarily with automation, and this is also true of healthcare. Pharmacies, assembly lines, and check-ins are all commonly automated systems. Although this appears a concern for humankind, it can facilitate more complex roles for humans, moving the species more from physical work to cognitive labor.

Healthcare organizations exploring the use of AI are understanding how to retrain or redeploy their staff, rather than make them redundant. As these AI systems become increasingly common and grow in capacity and accuracy, conversations are moving past manual and cognitive labor to how specialisms such as radiology will be affected. Healthcare teams that use image processing AI, for instance, are increasingly confronted with the dilemma of what to do with teams of specialists. AI can be used to redirect resources to amplify patient care, with more opportunities to engage with patients and their treatment. AI can be understood as an efficient and effective tool that can analyze and detect patterns in patient data at a rate and efficiency inconceivable to humans, with little misinterpretation. However, AI should be seen to augment a human's natural knowledge and intellect rather than be left to make a final decision.

Healthcare is unique in that a digital-only approach is inconceivable. It appears that human intervention will always be required to mitigate risk and confirm predictions. Even with the most sophisticated AI, 53% of 1,000 respondents to a Diabetes.co.uk survey stated they would want human verification of a machine-led diagnosis.[109]

Affecting the future

There are many indications that how people are using technology is changing how children and adults develop. Among many things, research now demonstrates too much screen time is associated with structural and functional changes in the brain, affecting attention, decision-making, and cognitive control.[110] Another study of 390,089 people by the University of Glasgow found an association between a high level of time in front of a digital display and poor health.[111] Screen consumption time is considered to be sedentary behavior. Excessive screen time has been demonstrated to impair the brain, particularly the frontal lobe, which is used by humans in all aspects of life. The next few decades will illustrate the implications of connected technology on human physiology, psychology, evolution, and ultimately survival.

Playing God

Engineers from the University of Adelaide developed an AI agent that was able to predict when you were going to die. The AI agent analyzed 16,000 image features to understand the signs of diseased organs.[112] The neural-network-based AI model was able to predict mortality within 5 years with a 69% accuracy rate. Enhancements in data access, AI, and predictive analytics is enabling humans to predict more precise morbidity and mortality risks, which reignites the debate as to whether AI is playing God or enabling humans (facilitated with a data-driven approach) to play God.

Overhype and Scaremongering

The capabilities of AI have been historically overhyped through marketing and simplified media representations. Media scaremongering, along with more public breaches of data, have also not helped the cause of AI. There is a fog of confusion and misconception, particularly over the possible use (and misuse) of data and outputs, which needs to be resolved for stakeholders to seriously consider implementing AI into healthcare systems.

Stakeholder Buy-In and Alignment

It is pivotal to ensure all stakeholders are aligned in the health intelligence a project can provide and how it will be used. For instance, developers of AI for healthcare applications may have values that are not always aligned with the values of clinicians. There may be the temptation, for example, to guide systems toward clinical actions that would improve quality metrics but not necessarily patient care, or skew data provided for public evaluation when being reviewed by potential regulators.

Policy, Law, and Regulation

The next decade will provide a pivotal juncture in the ethics of AI. Regulation in the form of laws, ethical guidelines, and industry policy will influence the direction of machine learning. So too is machine learning influenced in the absence of such policies and where much of the risk is considered to lie. By failing to set a perimeter of acceptability, we open the risk of wider unintended consequences.

Technical governance is needed, particularly for the application of AI. It is vital to engage key opinion leaders, stakeholders, lawmakers, and disruptors to form and nurture these policies to ensure the potential of machine learning is realized without undue restriction. Applications of

AI in non-healthcare sectors have demonstrated that there are potential ethical problems with algorithmic learning when deployed at scale.

Lawmakers and critical stakeholders in the legal decision-making process must understand the complexities of AI and welcome the challenge of the first-time problems it brings. Furthermore, it cannot be assumed that AI will be solely used for benevolent purposes: AI is already used in warfare, with the United States and China spending heavily.[113] (Interestingly, the UK has decided to focus on AI ethics.)[114] Healthcare professionals using AI must understand how deployed AI algorithms are created, assess the data used for modeling, and understand how the agent can safeguard for mistakes. Regulators must be prepared to handle each concern individually, as all AI is individual, akin to humans. Safe and inclusive AI development will only be achieved through a diverse, inclusive, and multidisciplinary team of engineers, humanitarians, and social scientists.

International engagement between organizations and governments on the critical topics raised by AI is a must. Questions of ethics, responsibility, employment, safety, and evolution are examples of conversations that will require stakeholders from the public sector, private sector, and academic sector to collaborate as AI embeds itself in society and to ensure humanity's democratized prosperity.

Data and Information Governance

Implementation of AI in healthcare requires addressing ethical challenges such as the potential for unethical or cheating algorithms, algorithms trained with incomplete or biased data, a lack of understanding of the limitations or extent of algorithms, and the effect of AI on the fundamental fiduciary relationship between physicians and patients.

A data governance policy is a documented set of guidelines for ensuring the proper management of an organization's data—digital or otherwise. Such guidelines can involve policies for business process

management (BPM) and enterprise risk planning (ERP), as well as security, data quality, and privacy. This differs from information governance, which is a documented set of guidelines for ensuring the proper management of an organization's information. An information governance policy would cover the use of predictive analytics outcomes. An ethical code of conduct or ethical governance policy would cover the organization's direction and intent and steer governance.

Is There Such a Thing as Too Much Policy?

There is a possibility of too much policy, hampering progress and disengaging stakeholders. AI has fueled innovation across all industries including digital health. Start-ups have led the way for AI pervasiveness because of less bureaucracy and red tape compared to organizations that are larger and typically slower to respond to environmental change. Audit your policies to ensure space for listening to the user's voice and incorporating it into cycles of ethical innovation.

Global standards and schemas

There are several bodies that exist that share common codes of ethics and responsibility, but no globally enforced schematics. An AI agent that predicts your risk of death has little validation other than you, or someone else it has predicted on, dying. How should such an agent be verified? To which standards must it uphold? Must the agent, or its creator, provide evidence? As yet, AI has no defined industry standards to which technologies are tested against. As such, it is difficult to be assured of the quality of the AI you are engaging with. And in a fast-paced industry such as AI, the concerns presented are novel and immediate. Global standards are required to set a minimum expectation of AI systems.

Healthcare also suffers from a time lag from evidence-based to clinical acceptance. For example, people with type 2 diabetes have placed their type 2 diabetes into remission for well over a decade, but health systems have been slow to be able to register a patient as placing their type 2 diabetes into remission. Thus, patients have historically either been labeled as having reversed their type 2 diabetes or as living with type 2 diabetes. If patients were labeled as having reversed their diabetes, they often would not receive the tests that confirmed their type 2 diabetes was in remission and were at risk of not receiving checks. Equally, if the patient was labeled as having type 2 diabetes (and did not), this has an impact on insurance and is technically untrue. Collaboration on schemas and parallel strides of adoption would help to reduce international disparities.

Do We Need to Treat AI with Humanity?

The basic concept of reinforcement is parallel to how humans and other animals learn. For instance, when training a dog, compliant and expected performance is often rewarded with a treat. Noncompliance by a reinforcement learning system is met with punishment. This is similar to reinforcement learning by an AI agent, where we build mechanisms of reward and punishment. Positive performance is reinforced with a virtual reward, and negative actions are penalized to augment avoidance.

The majority of AI systems are currently fairly simple and reductionist. As AI develops, we can expect them to become more complex and lifelike. Humans are already beginning to develop relationships with robots, in many cases seen as replacements for another human whether for companionship or carnal desire. It could be considered that a penalty input to an AI system is a harmful or negative input. Are genetic algorithms that delete generations that are no longer of use a form of murder? At what point do we need to consider AI's humane treatment is catalyzed when AI systems can mirror both human behavior and appearance?

If AI systems are considered to be able to perceive, feel, and respond to stimuli, it is not a huge leap to consider their place as a species, legal status, or even nationality. An AI agent made in America was even given Saudi Arabian citizenship.[115] Should AI be treated like an animal of intelligence comparable to humans? Can machines truly feel and suffer? The moral questions of how we treat AI agents, mitigate suffering, and the risk of negative outcomes are fundamental to responsibly progress AI toward its immense potential to better the lives of humanity.

Employing Data Ethics Within Your Organization

Many organizations have a code of conduct that serves as an internal communication on expected behaviors, requirements, and engagement and provides external stakeholders confidence and reassurance. Codes of conduct should be shared with employees alongside appropriate training to ensure the code of conduct is understood and observed. All stakeholders of an organization should practice and promote the code of conduct adherence.

Ethical Code

A code of ethics or ethical code is a document that is used to govern the moral conduct of an organization. A code of ethics demonstrates that an organization is committed to responsible business and technological progress.

The code of ethics narrates the behaviors that are promoted by an organization and those that are considered harmful to the organization's own moral compass, reputation, or clients. It may not cover actions that are illegal, but it will typically state the repercussions of noncompliance and how such violations can be reported. Employees should be made

aware of items that are not obvious to them and then facilitated to avoid inadvertent yet potentially harmful actions. A code of ethics should also include a summary of the motivations for using data and its purpose in the organization's ambition—and reflect the profitability, integrity, and reputation of a business.

A code of ethics should be free of technical and philosophical jargon and directly communicate the expectations required of employees. Keep it simple and right to the point. Set the expectations for new employees at the time of joining and develop an organizational culture of adherence. Ensure all stakeholders are aware of amendments and reviews. Rather than only referring to the code of conduct to recruits or clients, refer to the ethical code frequently to engrain the ethical direction of your organization among internal and external stakeholders. Refer to the code of ethics when examining the ethical risks of incorporating new technologies into the decision-making process.

When developing your code of ethical conduct from scratch, consult employees and stakeholders for their contributions. Questions to consider asking include the following:

- What does AI ethics mean to you?
- How can what we do improve humanity?
- How should our organization act responsibly in its ambition to [insert ambition here]?
- What are the potential benefits of what we would like to achieve?
- What are the potential disadvantages of what we would like to achieve?
- How can we improve our ethical code?
- Are there items in the code of ethics that are confusing or require more explanation?

- Is the code of ethics useful in making decisions?
- Are the organization's ethics in alignment with your own ethical viewpoints?

Once all of the preceding questions are answered by different segments of the organization, a consensus can be reached on an implementation plan. An organizational code of ethics serves as a reference point for disciplinary actions for those who fail to meet the standards. The code of ethics provides a solid foundation for identifying and dealing with ethical challenges.

Ethical Framework Considerations

Organizations have a responsibility to be ethically guided in the collection and use of patient data. Organizations require a data privacy approach that prevents breaches and enforces security. Failing to adhere to regulations can lead to fines, reputational repercussions, and the loss of customers. There are various ways to mitigate the risks of data collection while taking advantage of the opportunities big data has to offer. There are several techniques to safeguard ethical data collection.

Collect the Minimal Amount of Data

The first step toward protecting user data is collecting only the data that is required for the task at hand. More data does not always mean more useful data. Although machine learning typically welcomes more data, data collection must be kept deliberate and concise. For instance, there is little to no use in knowing someone's ethnicity and BMI if they are applying for a credit card; but this same data is important when determining an individual's risk of type 2 diabetes.

Identify and Scrub Sensitive Data

All sensitive data should be identified and secured. Data scientists should be aware of what data is deemed personal and how to use such information. For instance, data collected on consumers that is collected without their consent should be scrubbed of personally identifiable data. It is paramount that only individuals with the appropriate permissions can observe sensitive data. Many organizations also discourage the use of USB sticks and external data drives to ensure data remains safe and on-site.

Compliance with Applicable Laws and Regulations

Trust in your organization's approach to data, and information management can be garnered through adherence to appropriate local and national standards, policies, and laws. There are several data and information governance laws and regulations that set the boundaries of legal and appropriate data usage. It is vital that all regulations are abided to. Bodies such as the FDA and MHRA (Medicines and Healthcare products Regulatory Agency) set boundaries for the legal, ethical, and compliant collection, usage, and management of data. Compliance with GDPR and approval from governance-related accreditation bodies demonstrates your organization's approach to ethical data usage and builds trust with your users.

Many national- and international-level standards will have a fair degree of overlap, requiring appropriate data reporting, documentation, transparency, and risk mitigation procedures. Within Europe, for example, adherence to regulations such as GDPR is a necessity and regulates over 27 countries in the European Union. International standards such as ISO 27001 ensure good data and information management standards. ISO 27001 (formally known as ISO/IEC 27001:2005) is a specification for an information security management system (ISMS). An ISMS is a framework of policies and procedures that includes all legal, physical, and technical controls involved in an organization's information risk management processes.

Many data regulations are self-accredited. This means that as an organization, you must collect the documentation required to demonstrate adherence to standards without the requirement for it to ever be verified by an external body. The concerns of self-accreditation are apparent, particularly when applied in healthcare.

Robust data and information governance is the foundation of data ethics.

Thorough data and information governance is the foundation of data ethics. Governance is concerned with how data and information is protected in its collection, use, analysis, and disposal. Good governance covers a variety of topics including data architecture, infrastructure, risk assessments, audit reports, risk treatment plans, design workflows, information security policies, complaint procedures, medical governance, corporate responsibility, and disaster recovery planning. An organizational ethical code steers the direction of good governance.

A Hippocratic Oath for Data Scientists

Data scientists with access to sensitive data are typically required to sign documentation to ensure they comply with risk mitigation and security policies when they begin employment. It has been suggested that to engrain a culture of safety and positive progression in AI, data scientists acting in digital and nondigital sectors pledge allegiance to a Hippocratic oath to do no harm. This would sit alongside regulations, internal manifestos, standards, and codes of conduct operated by an organization.

Auditing Your Frameworks

Conduct regular auditing of ethical, data, and information governance procedures and rehearse worst-case scenarios to ensure that if the worst were to happen, your organization would be able to minimize as much risk in as little time as possible. Risk mitigation drills should be rehearsed at least twice a year.

Review each section of your code of conduct to ensure it represents the values of your organization. Technology organizations move rapidly, so ensure compliance with the latest standards, guidelines, and policies relevant to your industry. Organizations operating with medical data may come up with new methods of prediction, or new AI abilities. Consider aspects of your code of conduct that are missing, particularly if your organization is scaling, or facing new use cases. For example, a historically business-orientated organization that begins to engage with patients will require a new section of engaging with patients.

There is no reason not to engage every member of staff in developing and reviewing your code of conduct. Digital surveys are a quick and easy way to funnel opinion. Diverse contributions make for diverse, inclusive, and multidisciplinary approaches to AI. Ask for comments on new sections or concepts you feel employees should be aware of. Conversations can highlight areas in which employees may need further training, awareness, or demonstration. Develop a dialogue with employees to engrain the organizational culture within the wider team and foster an ethically aware environment. Having the maximum number of stakeholders on board with your organization's code of conduct will likely increase employee motivation and productivity.

An ethical framework in healthcare organizations, particularly for those using AI, provides internal and external stakeholders the security, transparency, trust, and direction required for maximizing engagement. Adhering to industry standards and involving members of your organization in developing and maintaining the ethical framework will accelerate employee and organization-wide adoption. Moreover, good policy documentation, training, and auditing help maintain the integrity of organizational AI.

CHAPTER 7

Future of Healthcare

“Medicine is at a critical juncture where it is necessary to keep up with the speed of discovery in the real-world.”

—Charlotte Summers

Connected healthcare and AI, when applied to healthcare, brings with it the ability to transform the lives of communities globally. People demand more today than ever before, and healthcare is expected to keep up in a golden era of innovation at the same time as a paradigm shift from healthcare volume to value—patient numbers to patient outcomes. As awareness of AI in healthcare grows, so too does public expectation that it will be used to improve day-to-day experiences.

New advances in medicine and concepts that were once the futuristic topics of science fiction are gradually becoming a reality. Gene therapies, 3-D printing of human organs, liquid biopsies, robot-assisted surgeries, and voice-enabled personal assistants are now realities that are becoming more sophisticated as time goes on. Advancements in technology are affecting not only practiced medicine but the public perceptions and attitudes toward health, lifestyle, and what it means to be healthy. Healthcare must implement innovation wisely to engage as many patients as possible. Digital health technologies evolve at a lightning pace. The impact of data science, AI, machine learning, and connected healthcare technologies is tremendous and requires an unbiased mind and willingness to engage in a continually evolving environment with as much knowledge as possible.

Shifting from Volume to Value

Healthcare itself is transitioning from a volume- to patient-centric value. There is a shift in medicine in which pharmacology is not the first and only line of treatment; and there is a greater emphasis on the role of lifestyle as medicine for preventative and therapeutic purposes.

Traditional payment systems in healthcare are grounded in rewarding providers for the volume of patients referred and treated. Volume-based care focuses on economies of scale. Providers receive a discrete amount of money for all of the services provided to patients for a set duration of time. Value for money, patient experience, and healthcare quality are subsequently secondary considerations for evaluation. Providers are incentivized by patient volume and the cost of the overall care provided to patients. Healthcare resource and cost pressures, exasperated doctors, and incentivized clinicians have unknowingly facilitated a “see as many patients as possible” approach that has driven disease management from a medication-first, patient-second perspective. Hospitals and clinicians are incentivized to see as many patients as possible, conduct as many tests as possible, and approach disease from a medication-first perspective.

As a result, volume-based care models are typically modeled on financial metrics such as optimizing profit or minimizing cost per patient rather than health outcomes.

Volume-based care payment systems often penalize providers financially for keeping people healthy, error reduction, or complication reduction, as they are not a driving KPI. For instance, if you were to be diagnosed with type 2 diabetes, your healthcare team would typically do the following:

- Be given information on how to manage blood glucose levels per current guidance.
- Progress on a treatment regime.
- Be given information about dietary choices.

In the UK, for example, if a doctor's surgery or clinic were to enable the patient to realize type 2 diabetes remission, the surgery or clinic would fail to receive the money it would have for the drugs no longer prescribed to the patient. Personal biases can also complicate matters. In the UK, for example, doctors involved in assessing which drugs should be prescribed to NHS patients were found by The Telegraph to have received £100,000 per year from pharmaceutical companies.[116]

Volume-based care is focused on a finite amount of money for a population. With populations growing and resources stretched, the evolution toward value-based care has been expedited. Healthcare is now in a transition toward patient-based care, which is synonymous with value, rather than volume.

Tseng et al. defined the value in healthcare to be the quality of care, typically measured through healthcare outcomes.[117] A value metric also for consideration is patient experience or patient-centeredness. Patient-centeredness is an important but not necessarily dominant quality measure for value-based care. Patient-centered care utilizes a multifaceted approach, revolving around the patient, their goals, and the wider family in decisions and evaluations. There are several domains to patient-centered care, which includes patient-centeredness.

Patient-centered care is a collaborative and multifaceted approach involving delivering health and social care focused on patient experience and patient outcomes. Patient-centered care focuses on empowering patients with the knowledge and experience, skills, and confidence to manage and optimize their health. Care is compassionate, personalized, coordinated around the patient, and respectful of the patient's worldview. These goals may appear commonsense for healthcare practitioners. However, they have not always been performed as standard practice. Healthcare has typically been reductionist in its approach, addressing a sum rather than the whole and acting for patients rather than with patients.

Organizations like Buurtzorg Neighbourhood Care in the Netherlands are an exemplars of patient-centered care. The organization delivers care to 80,000 patients a year with 10,000 nurses, 21 coaches, and two managers.[118] The collaborative approach focuses on empowering patients through relationships with caregivers. Over time, patients see improvements in health status while the total external care time required is reduced. Patients and families are enabled with the skills and confidence to look after themselves.

Although two distinct philosophies, patient-centered care and value-based care are becoming synonymous. Definitions of success incorporating patient-centered outcomes, perspectives, experiences, and preferences are positively influencing quality and delivery of care.

Value-based care does not mean cheap or economy care. Indeed, care is evaluated through the quality of patient experience and outcomes; and the model reshapes how care is received by patients.

Wellness and prevention are emphasized in value-based care. Lifestyle, behavior, and environmental factors account for 90% of disease risk.[119] This has been evidenced by the growing prevalence of type 2 diabetes across the world. With growing pressure on healthcare resources, lifestyle medicine has silently refocused efforts on the role of lifestyle as medicine. Evidence demonstrates that activity, nutrition, sleep, and stress all impact physical and emotional health markers.[120] Lifestyle medicine's focus on lifestyle and behavior enables healthcare professionals to deal with reversible, non-communicable diseases such as type 2 diabetes and hypertension. Rather than treating the patient in a reductionist approach, the patient is provided with holistic, compassionate care. The prescription of lifestyle medicine is already taking place in healthcare, with digital applications prescribed to treat disease. Type 2 diabetes patients, for example, may be prescribed the Low Carb Program app from Diabetes Digital Media, demonstrated to place type 2 diabetes into remission for 1 in 4 patients.[121] Lifestyle interventions are typically safe, easy to implement, and scalable. Growing levels of chronic disease can be

mitigated with a simple approach delivered in an engaging and effective setting.

Instead of reacting to events, there is a preventative, lifestyle-first approach to health in value-based care. The focus is on patient wellness, healthcare quality, and efficiency. This is where patient data, digital health, and AI hold a fountain of promise. The emphasis on efficiency ensures that the approach is scalable and effective, enabling providers to reduce healthcare costs and improve patient clinical health outcomes. Prevention of disease through quitting smoking, nutritional change, lifestyle change, activity, sleep, and identifying genetic risk factors reduces the burden on healthcare resources. Wellness is becoming incentivized. It has become the best interest of insurance companies and healthcare providers to monitor their patients' health more closely. Start-ups and digital technology companies are developing digital health tools that are disrupting the traditional doctor-patient relationship and have enabled providers to become third actors involved in facilitating sustained health. The goal of value-based care is to standardize healthcare processes through best practices and to democratize access and quality of care. Mining of data and historical evidence can determine which methods work and which don't.

Keeping people well reduces healthcare delivery costs and optimizes the use of resources. For instance, when managing a chronic condition such as type 2 diabetes, value-based care uses a collaborative and multidisciplinary approach to managing the disease to prevent diabetes-related complications. Patients engage with one healthcare team that is aware of patient progress and health status. The healthcare team may contain the patient's diabetes nurse, dietician, behavior health coach, and other professionals to support patient progress. The team would set patient-centered goals and help a patient with the following:

- Maintain blood glucose control.
- Introduce a digital community of people with type 2 diabetes for support.

- Provide health mentoring and habit maintenance support.
- Facilitate an activity program.
- Use the latest evidence base to provide nutritional guidance.
- Deal with the psychological aspects of type 2 diabetes.

Incentivization is transformed. Take a hospital setting, for instance, in which rather than being paid on how many patients could be seen, the hospital is being compensated based on the opposite. The hospital is paid based on how many patients are in positive health status and how many beds are available. The focus moves from the point of hospital admission to predicting the likelihood of future risks and anticipating them with preventative solutions. As well as face-to-face delivery, ongoing patient support can be provided digitally, for example, with health coaching, delivering an exercise program, or help with mental health concerns through apps, wearable technology, or telemedicine. AI and predictive analytics enable healthcare providers to optimize the delivery of healthcare, focusing on preventing and treating disease. Value-based care can utilize low-cost digital technologies to enhance and democratize care. The simultaneous collection of behavioral, demographic, health, and engagement data provides an opportunity for machine learning and development of novel AI, to rapidly improve, and learn from, user behavior and outcomes.

Managing healthcare operations and subsequently evaluating healthcare quality is complex. To assess a value-based care model, healthcare organizations must collect and analyze data, objectively assessing performance through the quality of care, patient health outcomes, and efficiency of cost. Providers can report and model preventative care metrics such as hospital readmission rates, error

rates, disease progression, population health improvement, and engagement strategies. Healthcare quality is defined by several metrics. Patient experience and satisfaction scoring is typically the first metric of evaluation. Metrics such as time spent per patient, patient engagement, medication savings, and adherence are examples of quality that relates to productivity.

Many applications of AI in healthcare are within a hospital setting. Hospital management platforms are delivering a return on investment by enabling organizations to prepare. Systems that predict times of peak patient traffic, predict readmission times, and use real-time data to cope with the demand of real-world healthcare are seen as long-term systems of value to healthcare providers. Applications involved in improving clinical care demonstrate enormous potential. Long-term studies are required to evaluate the impact on disease management and population health.

Evidence-Based Medicine

Evidence-based medicine is underpinned by decision-making that is based on the latest, most reliable scientific evidence. As shown in Figure 7-1, it is an approach to solve a clinical problem by integrating best research and clinical evidence with real-world clinical expertise and patient values. In the case of healthcare, this is deemed to be randomized controlled trials and increasingly real-world evidence. Datafication and the IoT are contributing significantly to the evidence base. David Sackett, a pioneer of evidence-based medicine, defined the concept as “the conscientious explicit and judicious use of current best evidence in making decisions about the care of individual patients.” [122]

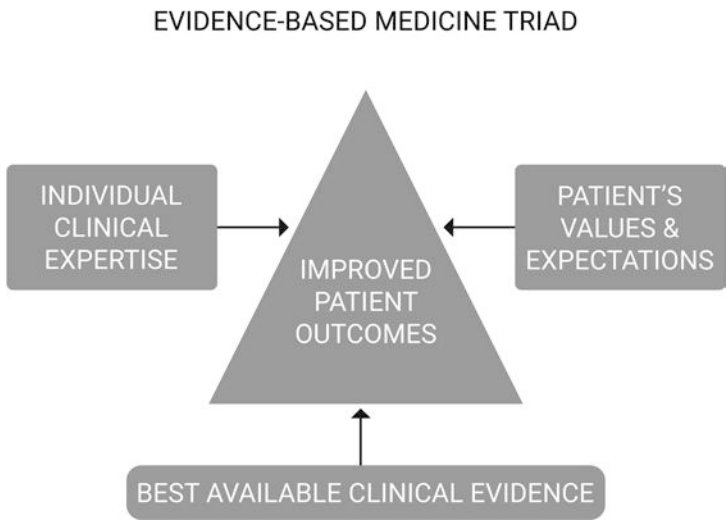


Figure 7-1. *Evidence-based medicine*

Interestingly, the traditional evidence base for pharmacology has been demonstrated to be biased, leading to questions over validity.[123] Randomized controlled trials (RCTs) are deemed as the most robust and reliable form of evidence for assessing the efficacy of a treatment. Evidence demonstrates that many factors can influence the reliability of RCTs, including methodological quality, reporting quality, and source of funding. Pharmaceutical companies fund the majority of clinical research that is undertaken on medications and face a conflict of interest. Studies have demonstrated bias in favor of industry in research funded by industry, which undermines confidence in medical knowledge.

Real-world evidence comes at a critical juncture for evidence-based medicine. Real-world evidence is disrupting traditional evidence-based hierarchies and approaches. Datafication of human experience through mobile phones, social media, digital communities, health apps, nutrition tracking, wearables, and health IoT has empowered patients to become their own evidence base and influence healthcare academia and understanding. Patients, for example, can compare the results

of different hospitals, or even individual doctors, to best choose their treatment providers.

Most RCTs conducted within pharmacology are internally valid, ensuring an exact population is involved in the testing of a particular drug. For example, a trial for type 2 diabetes medication may focus on people who have type 2 diabetes only and no other comorbidities; whereas in the real world, the comorbidities of hypertension and high cholesterol are often present in people with type 2 diabetes. This approach means there is often a gap between what is reported in RCTs and what is reported by patients. This is where digital technology can help bridge the gap between pharmacological intervention and real-world patient experience. Digital communities, for example, can provide real-world evidence on medication side effects. As an example, health community Diabetes.co.uk provides medication side effect data back to pharmaceutical organizations through a mobile application in the form of analysis of patient discussion and reported side effects. The world's second-most commonly prescribed drug for diabetes, Metformin, for instance, has a side effect of diarrhea reported to affect 1 in 10 people taking the drug. Real-world evidence from Diabetes.co.uk members demonstrates this is 400% more common and reported by over 48% of people with type 2 diabetes.[124]

Real-world evidence has historically been viewed as anecdote by the medical community. However, the datafication of patient lives now means that data is directly received from devices, wearables, and sensors, ensuring that data cannot be misreported. Real-world evidence is growing at a pace that healthcare is struggling to keep up with. Digital health platforms and the aggregation of data also calls into question previous paradigms of medicine. The Low Carb Program from Diabetes Digital Media, for instance, was able to aggregate data from 100,000 people to demonstrate that a low-carbohydrate diet could be used to place type 2 diabetes into remission. Type 2 diabetes was considered, 5 years ago, to be a chronic and progressive disease. Big, real-world evidence is proving that it doesn't have to be.

It is a necessity for evidence-based medicine to keep up with the developments of real-world evidence. As David Sackett said, “half of what you learn in medical school will be out of date in 5 years of your graduation; the trouble is nobody can tell you which half. The most important thing to learn is how to learn on your own.” Patients now use the Internet, themselves, and their peers as the evidence base; and ubiquity of real-world evidence and AI in healthcare means that healthcare professionals can no longer be ignorant to change in the evidence base and biased toward old paradigms.

This data is a wealth of opportunity for AI and machine learning. Healthcare can adopt a data-driven approach, utilizing a variety of data sources to improve the patient experience and reduce costs. The exploration of this data over the ensuing decades will significantly progress healthcare and facilitate precision medicine.

Personalized Medicine

Around 10% of disease risk is based on genetics.[125] Each person has a unique version of the human genome. Groups of particular patients may share common genome characteristics and thus the risk of disease. For instance, research demonstrates people of South Asian descent have a higher prevalence of type 2 diabetes compared to British Caucasians in the United Kingdom.[126] Personalized medicine, also known as stratified or precision medicine, is an approach that stratifies patients into groups and makes informed clinical decisions for the delivery of treatment and interventions based on the patients’ anticipated response. Personalized medicine approaches to disease management are tailored to the patient. Patient health is managed on an individual level to achieve the most optimal state of health possible. For example, our genetic variations determine how our bodies would respond to a particular drug. One drug might not fit everyone’s requirement. Two

people taking the same dose of the same drug might respond differently. Using personalized medicine, a right combination of a drug and its dose can be selected for everyone. This approach to medicine is not new, and healthcare professionals have been using this approach since the time of Hippocrates. However, the hyper-personalization achievable from the use of the patient genome and its falling cost—and medical data from health records, wearable technology, and the health IoT—has roused the public interest once again. It has never before been possible to predict the risk of disease, how the human body will respond to a particular medication, or print treatments out of materials other than ink. The combination of these technologies will fuel an era of personalized care and healthcare innovation.

Predictive tools can be used to evaluate health risks and develop personalized healthcare plans to mitigate patient health risks, prevent disease, manage disease, and treat a disease precisely if it occurs. As healthcare becomes increasingly personalized for patients in both treatment and service delivery, it is critical that access is widened to ensure participation from all groups of society.

Diagnostic tests, such as blood tests, are typically used to identify appropriate treatments based on a patient's physiological analysis. The selection of optimal therapies will be increasingly personalized to the patient genome. Healthcare providers will be able to diagnose current illnesses, predict future risks of diseases, and identify predicted response to treatments and subtle traits within an instant. Genetic testing has started to make an impact in personalized medicine; DNA results are limited to being imported into services that can personalize treatment regimes, diet plans, or education to a patient. Currently, DNA test results take weeks to receive. An era of instant genetic test results will empower patients to make informed choices about treatment, services, products, medications, and outcomes.

Raising serious ethical issues, the decision to have a genetic test and the implications of its results deserves careful preparation and thought.

For providers, it provides an opportunity to fully engage in personalized, value-based care, with the quality of patient experience paramount. For patients, it can pinpoint genetic mutations that predispose a person to disease. The physical and mental consequences of such knowledge can be profound and have caused controversy among the medical establishment and general public. A natural division among all people, regarding specific issues, is to be either for or against the topic of discussion. Progress is not without moral consequence, and legislation must endure the speed at which first-time problems are occurring. People are conflicted about having a lot of information about themselves.

Vision of the Future

Within the next decade, DNA profiling of the embryo will take place from within the mother's womb, which will create an immediate profile of the person's health status and risk of disease and enable the development of health and lifestyle treatment plans from the beginning of life. Ethical concerns in fertility will move to tackle making decisions based on an embryo's genetic status. Genetic profiling will also make it quicker to detect potential concerns and alter or delete potential genetic defects or characteristics that are disfavored. Predisposition to disease will be estimated, with healthy life care plans developed for patients to follow as a daily routine. Patients will constantly be monitored, with their data fed back to update their health records and focus on optimal wellness—driven through innovation in wearable technology, innovation in medicine, healthcare delivery, the IoT, and smart homes, smart cities, and smart communities.

Patient data aggregation will facilitate the understanding of baseline health and enable healthy people to be algorithmically visualized. Data monitoring and predictive analytics will instantly alert health professionals and patients should there be a deviation to the norm—alerting users

to poor health, possible illness, or nudging the user over an unhealthy lifestyle. One will be able to download a clinically validated app that would be able to detect, diagnose, and treat many diseases before attending a physical doctor's surgery. Sensors will become far less invasive—chips under the skin and smart tattoos that keep constant connectivity.

Visiting the doctor's practice will be very different. Your digital personal assistant, perhaps Siri or Alexa, may suggest you see the doctor based on your voice sounding ill, or at least different. Social media platforms, health communities, and phones will alert you to mental health problems based on what you type. No matter whether it's a hospital in Malibu or Manchester, your health record will be accessible and available on your phone and validated through a distributed ledger such as Blockchain, to which healthcare providers would continue to add to over time.

Advancements in robotics, automation, and digital health enable healthcare professionals to focus their time where needed most: with patients. Using a combination of data collected from the patient's everyday life, phone, wearables, health sensors, connected clothing—and combining that with clinical tests, scans, and check-ups—will allow much-improved analysis and monitoring and focus on health optimization. This is especially effective against non-communicable, chronic diseases like type 2 diabetes and tracking the progression of progressive conditions such as dementia. Both health providers and bill payers will incentivize patients to track symptoms and health markers to aid the treatment pathway.

Potential health concerns identified from patient data would trigger an alert to the patient's doctor for validation by a human after initial AI investigation and prioritization. As a patient, treatment will be personalized, based on health records, genetic analysis, and AI analysis of patient data. Medicines will also include 3-D printed treatments and digital interventions delivered through mobile apps. Digital interventions, measured by engagement, and health outcomes, will be hyper-personalized, based on demographics, behavior, health, goals, and preferences. Discreet wireless sensors will allow instant notifications of

change, while algorithms and AI models will be used to diagnose disease and treat patients. Earlier diagnosis of health issues will become more common through constant monitoring before they become more serious.

Care will be delivered through a hybrid of digital and face-to-face engagement. Virtual and augmented immersive experiences will reinforce and maintain behavior change. Drones will deliver your medication to wherever you are unless your autonomous vehicle takes you there first. You will be instantly alerted when data sources report an adverse effect of your medication, and you will instantly be offered a best-suited, tailored alternative. A variety of integrated healthcare disciplines will be involved in patient treatment to ensure a holistic and human approach to healthcare. Innovations will allow doctors, nurses, and wider healthcare professionals to be freed up to do the more human parts of the job. Robotic AI will be used to carry out more physical tasks such as moving patients around, creating sterile environments, performing blood tests, radiology assessments, and so on.

Patient concerns will be dealt in, in real time. Should a patient be showing symptoms of atrial fibrillation, for example, a doctor may record a patient's heartbeat on their tablet and upload it to a system that confirms or denies the doctor's concerns. If irregularities are spotted, the video is instantly sent to a cardiologist who can provide a diagnosis and begin a personalized treatment plan for the patient. Appointments and follow-ups will take hours or days rather than weeks or months. A network of connected care means several experts can look at a patient's concerns simultaneously and give second opinions.

Precision medicine is another emerging approach, which refers to modifying a medical treatment considering individual variability. For this measure, three additional data sets are added to existing traditional patient medical records: data about the patient's environmental exposures, their lifestyle, and their genomic data.[127] This information can assist the doctors in identifying which approaches, treatments, and preventions will be effective for which patients.

Connected Medicine

Wearables and the IoT hold the key to connected medicine. Data captured by sensors in these contraptions are playing an increasingly powerful role in healthcare, facilitating the development of patient-centric healthcare systems. Many factors are accelerating the acceptance of wearable healthcare solutions, particularly their use in clinical trials and academic studies to monitor patient health and lifestyle factors. For example, studies can record participants' health vitals using an Android Watch, Apple Watch, Garmin, Fitbit, or another smartwatch device. Participants use apps to record their lifestyle habits, nutrition, activity, and medication adherence and to monitor medication side effects among others. Wearables and patient data are beginning to be used by insurance companies to incentivize wellness. Insurance companies have historically targeted such insurance products toward digital-savvy buyers, offering the latest gadgets to incentivize health improvements. However, incentivized insurance products will become commonplace for wider populations, proving particularly effective in the combat against non-communicable diseases. The delivery of connected, digital healthcare opens an opportunity to monitor, manage, and reverse disease, broadening the risk portfolio and creating a longer-living, reduced morbidity population.

As sensors become faster, smaller, and more capable, patient health records profiles will eventually comprise of detailed sleep analyses, details of continuous blood glucose monitoring, heart rate, blood pressure, and approximate calorie burn. A smartwatch will combine a variety of diagnostic tools, able to monitor blood pressure, heart rate variability, blood glucose, ketones, and more. Health sensors will become embeddable, biodegradable, and constantly connected, playing critical roles in such tasks as patient care. It is worth noting that evidence demonstrates fitness trackers have a limited engagement lifetime and do not improve patient weight loss.

Table 7-1 briefs some examples of connected medicine devices, explaining what device can be used where and why it can be useful. For example, immersive technology devices can be worn on the head for education purposes. Some of the wearables and embedded sensors include the following:

- Stress bands

Stress tracking and calming of the mind are two key markets for wearables. Wearables currently monitor breathing and heart rate, detecting signs of tension to improve breathing and the ability to reach a calmer state of mind. Fitness trackers offer mindfulness functions, and there are headbands that deliver waveforms to the brain.

- UV sensor

Northwestern University has developed UV sensors that precisely monitor a person's UV light exposure. The wearable sensors, which are small enough to fit on a human fingernail, are wafer-thin. As a known and potent carcinogen, this has the potential to reduce overexposure to UV, complications such as heatstroke, and melanoma.[128]

- Smart tattoos

Smart tattoos place the sensor in the skin with significant strides made by researchers at MIT and Harvard.[129] Smart tattoo ink reacts with the biochemical composition of the interstitial fluid to indicate the status. Developments will allow for

tattoos to degrade, lasting for as long as required and even only displaying in particular lights. For instance, high blood glucose could turn a blood glucose smart tattoo red, whereas low blood glucose may turn the tattoo blue.

- Smart medication

There are many directions to which medication is being made smart. Smart medication can tell a patient and their respective healthcare team as to the amount of medicine consumed, the time it was taken, and provide timely reminders if it looks like it may be missed. Today, Bluetooth-connected bottle caps and pill packets provide reminders to patients to encourage medication adherence. Bluetooth insulin pens and pen caps that log the time, amount, and kind of insulin injected typically push data from the device to cloud, with users engaging in a digital app interface to query their data. Smart asthma inhalers, for instance, can identify an oncoming attack before the wearer recognizes the symptoms.

- Smart insulin

A form of smart medicine, smart insulin is next-generation insulin that responds automatically to changing blood glucose levels. The lower or higher blood sugar levels are, less or more insulin is released, respectively. In 2015, University of North Carolina researchers reported that their

smart insulin patch, which rests on the outside of the body, will use a system of micro-needles to automatically detect high blood glucose levels and administer insulin through live beta cells appropriately. Because the beta cells are kept within the patch on the outside of the body, there is zero danger of them being rejected by the immune systems of people with type 1 diabetes.[130]

Patients immediately turn to Dr. Google at the first sign or symptom of potential illness. Patients are best placed to understand what is happening to their body, now visiting the doctor more prepared than ever. More informed, more aware, and more health-conscious patients now attend appointments, with their engagement used to personalize their healthcare plans. However, this also causes varying degrees of mental anguish and hypochondria. Patients must not forget that qualified healthcare professionals provide a much more complete assessment and explanation of health, particularly if up-to-date with the evidence base. This is perhaps a contributing factor for why the medical profession has not always welcomed a patient-centric approach. A 2018 Diabetes.co.uk survey reported 87% of patients had a better relationship with their healthcare professionals as the result of using the platform, with 75% reporting a better understanding of their condition. When this survey was conducted in 2015, the same number reported a better understanding of their condition, but only 20% of patients reported improved relationships with their healthcare team.[131]

Table 7-1. *Applications of wearable technology*

What?	Where?	Why?
Immersive technology	Head	Education
Military apparel		Behavior change
Helmets		Intelligence to intelligence
Mixed reality		communication
Smart contact lenses	Eyes	Blood glucose levels
Trackers		
Hearing aids	Ears	Sound
Headphones		
Trackers		
Odor detection	Nose	Smell
Smart tattoos	Arms/Wrist	Blood glucose
Trackers		Blood pressure
Patches		Oxygen saturation
Implantables		Ketone levels
Smartwatch		Education
Trackers		Rehabilitation
Clothing	Body	Rehabilitation
Chest straps		
Implantables		
Trackers		
Exoskeleton	Legs	Protection
Clothing		Rehabilitation
Embedded footwear		Health metrics
		Posture correction
		Rehabilitation

Connected medicine provides a wealth of opportunity in solving novel health problems, particularly with disease management and elderly care.

In 2015, the United Nations approximated 1.2 in 10 people were over the age of 65. Projections estimate that this will jump to 22% by 2050.[132] Healthcare expenditure on the elderly is a mounting concern, as it accounts for a higher share of healthcare expenditure compared to other age groups. Equally, non-communicable diseases such as type 2 diabetes and obesity are pandemics, with the WHO estimating 600 million people will develop type 2 diabetes by 2050.[133] About 1 in 4 people suffer from two or more chronic conditions.[134] With rising pressure on governments, payers, and manufacturers to reduce healthcare costs, elderly and long-term care require solutions that are prepared for the impending rise in numbers. The Internet of Medical Things (medical IoT) demonstrates a tremendous potential to help.

Disease and Condition Management

Non-communicable diseases including type 2 diabetes, hypertension, and cardiovascular disease can be monitored by both the patient and healthcare team. Heart monitors can detect heart rate variability and alert the healthcare professional in real time. The same concept is liberating for parents of children with type 1 diabetes using smart blood glucose monitors. Blood glucose measurements are sent directly to the cloud where parents can access their children's blood glucose in near real time and be alerted to potential hypoglycemic events. Digital apps can be delivered to reverse type 2 diabetes, manage epilepsy, and improve quality of life.

Novel utilization of technology can be used to minimize the risk of health concerns and their subsequent healthcare costs. Elderly care can be delivered through digital apps. Risks such as falling, for instance, can be mitigated through fall detection and emergency assistance. Accelerometers found in phones or watches can be used to detect falls or seizures, triggering notifications to carers of an emergency, its details,

and location as they emerge. Future technology could also contain safety equipment. ActiveProtective, a company from the United States, provide a care belt that detects and deploys airbags if it detects a fall.[135] Technologies such as this can help save on avoidable elderly-care and fall-rated healthcare costs.

However, the true value of the technology, especially connected wearables in elderly care, is realized in the machine learning and predictive analytics observed from the data collected. For instance, deviations from daily routines could indicate cause for concern and emerging physical or mental health concerns.

Virtual Assistants

Virtual assistants are not just entertaining; they can accompany those who live alone and provide support to the aging and elderly. Virtual assistants can empower people by answering questions, teaching new skills, setting verbal reminders to perform tasks such as reminding to take a medication and even controlling the home by performing tasks such as picking up the telephone or making a call. As homes become more connected, virtual assistants will assist people in every aspect of life and interface with all aspects of the home. Virtual assistants could call the emergency services on behalf of a patient, citing exact details and location. Connected cameras can track movement to detect falling, triggering the virtual assistant to reach out for help. Not confined to the realms of speech, virtual assistants in the form of robots can assist in particular situations within elderly care, performing roles such as helping people get of bed, the bath, or into a wheelchair. AI within such virtual assistants could even learn patterns, for example, when people may want to visit the bathroom. Interaction through voice and touchscreen enables communication and will increasingly empower aging patients with greater autonomy and reduce the burden on elderly care.

As the capabilities of virtual assistants improve, so too will their ability to determine the emotions of their master, whether it be happy or sad,

excited or depressed. Virtual assistants will interpret the syntax, semantics, and tone of the conversation to check for signs of mental or emotional health concerns.

Similarly, there is an application for virtual assistants in products to assist people with disabilities. Apps can be used not only to control lights, music, and heating but also for lifestyle management. Nominet, a UK Internet organization, developed an open source prototype called PIPs to cater to people with sensory or cognitive impairments.[136] People are empowered with audio and visual prompts, nudges, and reminders to go through their daily routines and to develop the skills and confidence to navigate their environment. The prototype is built with a low energy Bluetooth controller that sends and receives notifications between devices. The devices can be used to guide people through tasks such as washing their face, taking a bath, or taking their medication.

Remote Monitoring

Connected devices can provide support for caregivers and healthcare professionals to keep people out of a hospital and reduce their visits to the doctor. Apps and virtual assistants can provide medication reminders; fitness trackers will monitor user health and movement, alerting care providers if deviations in a routine were to occur. Connected devices can alert family and care providers of events such as using a fridge, having a bath, or opening the front door. Aspects of routine demonstrated in the data would be learned and respective accounts notified if there were to be a deviation from the norm. Emergency buttons can alert the relevant services, ensuring that the elderly or those with long-term conditions can receive medical intervention when required most. The combination of connected devices, apps, and virtual assistants make the home smarter and more conducive to mitigating risk, ensuring the healthcare provider/clinic is up-to-date with the latest overview of patient health while empowering the patient.

Medication Adherence

Forgetting to take medicine is easy enough, but the implications can be a problem. More than 50% of prescribed medications are not taken as directed.[137] Missed doses can result in slower recovery time, exasperate illness, or more serious consequences. A study by Benjamin et al. showed that getting patients to take their medications appropriately could prevent approximately 125,000 deaths per year in the United States alone.[138] There are many apps and digital services that remind patients when to take their medication, but medicine is going beyond even that. Smart pills with ingestible sensors seek to help doctors improve clinical health outcomes through notifying their respective app on confirmation that a medication has been consumed.[139] The notification is triggered when the sensor comes into contact with the stomach fluid. This enables doctors to measure treatment effectiveness and optimize the patient pathway.

Accessible Diagnostic Tests

People with long-term conditions and the aging typically need increased diagnostic tests. Instead of turning up to a clinic or pharmacy to have urine and blood tests conducted, smart sensors in portable devices allow the test to be conducted in the comfort of the user's home. Results can then be wirelessly shared to their healthcare team. Tests for cholesterol, cardiac function, HbA1c, fasting blood glucose, vitamin D levels, and insulin improve the convenience of performing repeated diagnostic tests, enabling quicker treatment, reduced risk of complications, and reduction in avoidable healthcare spending. As the database and evidence base expands, so too do the application of wearables in detecting signs of disease. Type 2 diabetes, for example, can be predicted from heart rate variability collected through an Apple Watch.[140]

Smart Implantables

Implantable devices that are connected to a smartphone, healthcare provider, and wider networks will be used to monitor and treat disease in real time. The possibilities from this data are huge, as wearables are external implantables that have the potential to collect metrics on biological markers from organs and tissue.

Digital Health and Therapeutics

Telemedicine, which can be traced back to the 1900s, refers to the clinical application of technology that uses electronic technologies to support long-distance patient care, patient and healthcare professional education, and health administration. Digital therapeutics, or digital health, is an enhanced form of telemedicine that brings together digital and genomic technologies with health, lifestyle, and human factors to deliver personalized medicine to patients, enhancing the efficiency of healthcare delivery.

There are more than 150,000 apps focused on health and wellness in the Apple App Store, which have been downloaded by over 50 million people.[141] Digital health tools allow patients to take a proactive approach to their health and wellness, looking to influence aspects of human behavior for the purposes of improving health. As more of our healthcare interventions are focused on chronic rather than acute diseases, behavioral therapy is surpassing pharmaceutical intervention as first-line therapy. Digital health and wider telemedicine have applications in a plethora of health conditions and are enabling healthcare and treatments to change rapidly. For instance, patients scheduled for bariatric surgery can be prescribed a digital health app to assist in weight loss in preparation for surgery; a patient diagnosed with asthma can be prescribed a connected asthma inhaler and app; patients diagnosed with epilepsy can be prescribed an app to manage seizures.[142] Consultations through Skype

of video-conferencing facilities, EHRs remote monitoring, digital health education, and transmission of scans or images among other applications, are all considered part of telemedicine.[143]

The world is truly mobile, and digital tools aim to deal primarily with minor primary care issues through electronic means. More serious issues would be escalated to the relevant healthcare professional, provider, or team. Digital health enables healthcare professionals to anticipate healthcare concerns rather than waiting for symptoms to present.

Measures of performance for digital health will be based on how convenient the experience was for the patient including patient health outcomes, quality of care, and the number of users using tools over the long-term.

Education

Digital health provides a perfect opportunity for education. Less than 1 in 10 patients with diabetes in the United Kingdom attends structured education within 12 months of diagnosis.[144] Research demonstrates offline education benefits typically decline 1 to 3 months after interventions cease, suggesting that learned behaviors change over time and further support is required to ensure lifestyle changes can be maintained. As a result, engagement with offline diabetes education courses is limited. Digital education is engaging by its very nature; it is available at any time and on any device at the patient's convenience. Digital education can additionally be personalized to users through the data exchange between user and provider with the presentation, and analysis of real-time patient data subsequently encourages personalized, goal-focused practice and sustainable behavior change. Education can be made patient-centered and person-led rather than the traditional teacher-learner environment. Digital therapeutics provide a substantial opportunity to empower patients with the knowledge, skills, and resources required to manage and optimize their health.

Lifestyle changes such as stopping smoking, moving more, and losing weight can all be tackled through behavioral interventions, simultaneously targeting the major chronic diseases. Chronic diseases and their management are behavior mediated. Digital health tools enable treatment to be personalized, democratized, and scaled. Real-world evidence has largely demonstrated the efficacy of digital tools. More evidence and randomized clinical trials are required to demonstrate causation. One area that is still unclear is who pays for digital therapeutics. If a doctor should prescribe a digital therapeutic, who should pay? The insurance company? The healthcare provider? The employer? The patient?

Regardless, digital innovation is revolutionizing disease management. Patients are increasingly demanding autonomy; wishing to be educated about their health, aware of the latest evidence base and engaged with their health.

Incentivized Wellness

Ill health is expensive. To employers and the economy, physical and mental health issues cost the United Kingdom more than £77 billion.[145] Employees lose an average of 30 working days each year due to illness or underperformance as the result of poor health. The economic impact globally chronic conditions such as cancer, type 2 diabetes, and mental illness could reach over \$47 trillion by 2030 according to the World Economic Forum.[146] Individual lifestyle choices, physical and mental health, all impact work performance. There is a significant challenge mounting for providers such as insurers and employers in improving absenteeism and presenteeism. To do this, providers are realizing the importance of holistic health—comprising physical health and mental health, social and human interaction, both in and out of the workplace.

As healthcare interventions focus on chronic conditions rather than acute illnesses, behavioral therapy is typically becoming first-line therapy over pharmacological intervention.

Wellness programs provide a convenient mechanism by which employers can support health and wellness in the workplace. Previously, incentivized wellness meant encouraging individuals to visit the gym and take more steps—by taking the stairs or walking to work. Employers often promoted healthy behaviors associated with employee well-being and reduced healthcare costs. Today, wellness programs and digital tools go beyond signposting and screening. Immersive, engaging programs that are remotely accessed empower individuals to modify their behavior with the goal of achieving positive clinical outcomes, and better health. Behavior change in the form of therapy, behavior coaching, biomedical data feedback, and sensors enable wellness to be truly patient-centered.

AI

AI has been widely adopted in critical industrial roles for decades and has finally started to take a leading role in healthcare. Connected medicine has provided fertile ground for machine learning models to be developed. AI seeks to place the doctor in our pocket. Among other things, AI has applications in disease prediction, cost reduction, efficiency improvements, automating manual tasks, and promoting us to alter our health. As data volume and variety increases, the capabilities of AI models will become more precise, more probing, and also more contentious. Life expectancy from birth has stopped increasing.

Mining Records

Mining of healthcare information systems holds tremendous potential for medicine. Systems hold considerable information about patients and their health, prescriptions, doctor's notes, and more. Healthcare data within such systems can be used to improve the quality of healthcare, reduce costs, mitigate mistakes, and improve and democratize healthcare quality.

However, knowledge discovery from the data contained in such systems is currently challenging due to variations in complexity, vocabulary, and standardization. Data mining provides an opportunity to extract relevant information from both textual and image-based archives. Mining also enables the discovery of patterns from data that can be used to build predictive models.

As discussed in Chapters 3 and 4, unsupervised learning is known for feature extraction, whereas supervised learning is suited to predictive modeling. Mining of records is beneficial for both the patient and provider. Mining of records can be used to identify high-risk patients or those with chronic conditions for whom a personalized intervention could be delivered. Record mining enables providers to find best practices and treatment to reduce claims and hospital admissions. Through comparison of symptoms, treatments, and positive and adverse effects, healthcare providers can analyze the pathways that improve outcomes for patient cohorts, enabling clinical best practice and standards of care.

Conversational AI

Conversational AI refers to systems that can talk. Rather than a user interface based on text or code input, individuals can interact with conversational AI system with their voice. Users are increasingly using chatbots to communicate with products and services. Voice-driven AI such as Amazon's Alexa can synthesize natural language to provide recipes or exercise tips, order products, or call a cab. The technology is catching on: 1 in 5 Americans owns a smart speaker, and there are over 100,000 Facebook messenger chatbots.[147]

As AI chatbots develop, so too are their capabilities. Intelligent personal assistants will act as healthcare assistants. By virtue of being voice-controlled, there are immediate applications for those less-abled, where only the voice needs to be used to perform tasks or instructions. Within health, conversational AI enables simple questions that do not

need the attention of a doctor to be answered. For instance, new parents could bombard a conversational AI with questions without fear of embarrassment or consuming healthcare professional's time. Questions such as what temperature a baby should bathe in, how often a baby should sleep, or whether there are developmental milestones taking place can all be instantly answered by an AI speaker. Many individuals turn to search engines to find the answers to their questions. However, most patients who query symptoms are unaware how to discern research quality and may come across conflicting evidence in addition to misleading information (or fake news) that can leave patients confused. With the assistance of a medical AI chatbot, patients can receive immediate assistance. Continuing the preceding example, if a child's new parents had a medical question or were concerned by a symptom (say, a chesty cough), it would be burdensome for them to visit the doctor for a response to every question. However, there is a need for medical confirmation. This cannot currently be detected through algorithmic means; and hence, the evolution of conversational AI in healthcare will be digital health assistants supported by and learning from healthcare teams and the world around them.

As conversational AI develops, cognitive systems will analyze conversation to detect early signs of mental, physical, or neurological illness. Voice-enabled devices such as Alexa will one day be able to identify symptoms of Asperger's, anxiety, psychosis, schizophrenia, and depression from conversational tones. This will assist doctors to predict better and monitor and track disease.

Imagine how much time and resources will be saved when virtual assistants, healthcare chatbots, and digital tools give answers to basic medical questions that do not require the intervention of a medical professional.

Making Better Doctors

AI will transform what it means to be a doctor. AI's integration into medicine is making better doctors and saving lives. From waiting times to prioritization, finding of evidence to maintaining productivity,

or supporting decisions—AI will assist doctors and healthcare professionals in making informed decisions.

Optimization

Suboptimal processes waste time and resources. Scheduling systems for clinics, doctors, and patients enable healthcare to be as efficient as possible. An AI system that was to schedule patient appointments and streamline communication would lighten doctor burden and focus time on more pressing matters. If a patient presented an urgent matter, for instance, an AI could evaluate and prioritize patient appointments.

AI will enable healthcare professionals to stay up to date with the latest evidence base. PubMed, for instance, has over 23 million papers within its archives.[148] AI that was to mine such articles and present the most relevant evidence would facilitate more up-to-date healthcare professionals and enable professionals to practice medicine to the best of their knowledge and evidence base. Similarly, AI can assist healthcare professionals in administrative tasks, performing roles of cognitive assistance.

AI can help mitigate a plethora of mistakes, including those made by humans. Record mining, in particular, will enable doctors to improve over time. For instance, digitalized data regarding patients, their treatment, healthcare professionals, and clinics can be mined to identify errors in treating certain conditions and avoid unnecessary hospital admissions. Zorgprisma Publiek, a Dutch organization, currently does precisely this: using IBM Watson to mine data to identify repetitive mistakes.[149]

Diagnosing disease

Professor Geoffrey Hinton, a pioneer of neural networks, famously stated that it is “quite obvious that we should stop training radiologists” due to the sophistication of image perception algorithms, which could soon

be more advanced than humans. AI can currently be used to examine medical scans, identify symptoms of type 2 diabetes, identify retinopathy and cardiovascular disease risk, and spot signs of breast cancer. AI will perform better as algorithms are trained on more data. However, it is not just clinical data that is required for disease diagnosis. Cardiogram, an organization based in the United States, can detect type 2 diabetes with an 85% accuracy rate—from Apple or Android watches, Fitbits, and other heart-rate enabled wearables.[150] Similarly, the US company Striiv can detect atrial fibrillation from 2 weeks' worth of heart-rate data. Algorithms to detect disease are showing great promise, with several already in medical practice. However, that does not mean all AI diagnostics are ready for implementation. Many AI tools are developed without peer review and its academic rigor. Essential details require verification, such as the algorithm's code, training, and validation datasets; the data to which it is compared; how performance is evaluated; and how neural networks come to their conclusions. As AI evolves, one would expect AI diagnostic tools to undergo the same RCT analysis that drugs are subject to. Diagnostic AI can save lives, money, and time through earlier diagnosis of conditions and responsive treatment.

Making and rationalizing decisions

Doctors make challenging decisions daily. It is vital that these decisions are as informed as possible. The use of AI, data mining, and predictive analytics is enabling clinicians to rationalize decisions and develop evidence-based treatment options. AI does not have to make the decision but can present the most reasonable opportunities to proceed.

Drug discovery

Developing a new drug is an expensive business, and only one in three new drugs will make it to the market. Every new drug brought to the market costs pharma an average of \$2.7 billion.[151] The consequence of

failed drug trials can be damaging: from reduced share price to closure of worksites and reduced staffing. Hence, the pharmaceutical and life science industry is increasingly turning to AI to facilitate drug research and development. Pharma isn't looking to replace humans. Instead, it is looking to see why the rate of failure is so high through using innovative AI. AI systems are being developed that can identify new therapies from information on existing medications; this could improve efficiency, drug development success, and accelerate the route to market for new drugs. In cases of pandemics, such as Ebola or Swine Flu, an accelerated route could save numerous lives.

There are approximately 10^{60} compounds that have drug-like characteristics.[152] The chemist of the future will be empowered through leveraging data: historical data, experimental data, inferences, and trends. AI enables the earlier identification of promising drugs and biomarkers. AI also has application in clinical trials. Clinical trial patient stratification allows the identification of patient groups most likely to benefit from the associated drug. This enables clinical trials operators to find the right patients for trials rather than patients who are unlikely to respond.

3-D printing

Innovation in 3-D printing promises exciting advances in the field of healthcare, where 3-D printing is beginning to disrupt traditional paradigms as the technology becomes more affordable and accessible. As the name suggests, 3-D printing involves printing 3-dimensional objects from a digital model, using additive processes to build the object. Also known as additive manufacturing, 3-D printing adds iterative layers, one on top of the other, to assemble the model. This allows precise modeling and reduces error. Although 3-D printing solutions are still in infancy compared to other technologies, the potential for its applications is exciting. Many AI experts consider it feasible that in the future, humans will adopt robotics and bioprinted material into their bodies. Soon humans

may be able to print just about everything, from prosthetics and pills to bioengineered replacement body parts and organs. The technological advance of 3-D printing has life-changing implications for patients.

Personalized prosthetics

The World Health Organization state that approximately 30 million people require prosthetic limbs, braces, or mobility tools and only 20% have them.[153] 3-D printing enables prosthetics to be custom built for each person. Through capturing a patient's measurements, prosthetics can be tailored to ensure comfortable, custom-fitted devices. Prosthetics will also become more comfortable as multi-material, 3-D printing assists with ensuring a better integration with the human body. Casts could also be made more comfortable through the same, personalized approach. 3-D printing facilitates the swift and cost-efficient creation of personalized products. Innovation in 3-D printing is combining integrated sensors and machine learning algorithms to support more natural, fluid movements. Predictive movement algorithms will evolve to mimic more natural movement, and humans will be able to use their brain and body to control them. 3-D printing technology enables prosthetics to be modeled and constructed in a little under 24 hours; importantly, printed prosthetics cost a fraction of conventional prosthetics.

Bioprinting and tissue engineering

Kidney, liver, and heart transplants could soon be a thing of the past. 3-D bioprinting marks the beginning of an era where transplant lists are a thing of the past. 3-D-printed organs can be built using the same techniques as 3-D printing, but instead use stem cells as the printing material. As well as printing cells, bioprinters typically output a gel to protect cells during printing. In the future, once these organoids are printed, they will be able to grow inside the bodies of patients. Princeton University has developed a bionic ear from 3-D printing tools, which can hear frequencies beyond the

human range. Printable and customizable implants from a patient's cells reduce the potential risk of rejection from the body.

One use case for bioprinting is the 3-D printing of human skin. Burn and acid victims, for example, have limited options for their disfigured skin. A team of Spanish researchers developed a 3D bioprinter that was able to produce human skin cells. Within 30 minutes, a biological ink containing human plasma and other materials was able to print 100 square centimeters of human skin.[154] The opportunities bioprinting technology enables are truly life-changing.

There is no doubt a cosmetic market for bioprinting. It is feasible that in the future, face printers could be used by people to apply a model of someone else's face to themselves, or even store a digital model of their face and regularly reprint their face for perpetual youth.

Alongside nanotechnology and genetic engineering, bioprinting may prove a tool in the pursuit of extending life spans.

Pharmacology and devices

3-D printing disrupts traditional pharmacology and device manufacturer approaches. For instance, the efficacy of drugs and other treatments can be tested on replicated human cell tissue. Through this, 3-D printing enables the printing of medications that are personalized for the patient. By reducing drug variability with this approach, there is potential for 3-D printed drugs to increase medication efficacy, reduce adverse effects, and improve adherence. The same concept applies to device consumables. Blood glucose test strips, for instance, could be 3-D printed for use by diabetic patients at home. The possibilities are endless.

Education

Alder Hey Children's Hospital in Liverpool, England, uses 3-D printing for teaching and familiarizing stakeholders with heart surgeries that are to be performed.[155] Every cardiology patient is unique; and for those

undergoing operations, 3-D models are printed to enable a variety of uses. Firstly, it enables surgeons to demonstrate surgeries that will be performed to often anxious and concerned parents using textures and colors that closely represent human tissues. Secondly, the 3-D models are used by staff and surgeons to familiarize themselves with the most important part of an operation—the child’s heart. Finally, 3-D models are used to train students.

3-D printing also enables surgeons to prepare before surgery through evaluating a precise model of the patient’s body or organs. This mitigates the likelihood of errors, improves surgeon accuracy, and reduces the amount of time the patient is required to be on the operating table. 3-D modeling is also useful in modeling bodily organs without the need for invasive procedures.

Gene therapy

Just like organ transplant lists, genetic diseases could soon be a thing of the past. The use of nanotechnology, which refers to working with tiny particles to manipulate cells and alter DNA, enables the editing of gene expression at a cellular level. For instance, gene editing has been successfully used to modify human immune cells to resist HIV infection. With 1 in 25 children born with a genetic condition, gene editing has the potential to treat or eliminate disease.[156] It could be used to correct defective genes in embryos and personalize the child’s immediate healthcare treatment plan. It is feasible that genetic conditions like cystic fibrosis, sickle cell anemia, and muscular dystrophy could be treated through editing the DNA in patient cells. Gene editing is not without ethical implications: many people have moral and religious concerns about the use of human embryos for research or the editing of genes. Also, there is a concern that nanotechnology and gene editing will only be available to those that are wealthy, which will exasperate the disparity in healthcare and interventions. Nanotechnology will significantly improve treatment and recovery time, with the absolute potential to prevent humans from developing disease at all.

Virtual and Augmented Reality

In the not too distance future, surgeons performing CT scans will be able to layer their scans over a patient's body via augmented reality; medical students will use virtual reality to explore the inside of the heart and burn victims will be virtually transported to a snow-covered mountaintop as a form of pain relief therapy. Immersion into a completely digital environment, virtual reality, has mainly been used for games. Innovation in augmented and merged reality means that we can now entwine the virtual and physical worlds and manipulate both environments simultaneously. Virtual, augmented, and mixed realities are increasingly being implemented in a wide range of medical applications. It is already evident that as it is incorporated further into healthcare, it has the potential to change the way many healthcare services are delivered.

Virtual Reality

Virtual reality (VR) is typically associated with gaming. In virtual reality, the user's reality is replaced by an immersive, entirely digital environment. This is currently achieved through a headset and handheld sensors, which enable interactions within the environment.

Augmented Reality

Augmented reality (AR) overlays a digital or 3-D environment in the form of objects, video, or data into the user's environment. The user's real-world experience remains central to the user's experience, and the information is added to the user's existing reality to enhance their experience. Information is distinctly digital and does not seek to emulate real-world objects. AR does not necessarily require additional hardware and can be

achieved through technology such as the mobile phone. AR's popularity was confirmed by the Pokemon Go app phenomenon in 2016. The game, downloaded more than 800 million times, overlaid 3-D characters within the user's real-world environment through the user's GPS location and inspired a wave of AR innovation.

Merged Reality

Merged reality seeks to emulate digital objects that can be interacted with. This requires additional technology such as a headset. Separate sensors track hand gestures and movement. In a mixed reality, the user can manipulate both the real-world and digital environment.

Pain Management

Research demonstrates that virtual reality environments can be used to reduce the amount of pain a human feels versus a control distraction condition.[157] The somatosensory cortex and insula, found in the brain, are linked to pain. Hence, through immersing patients in a virtual reality, VR will be used to enable patients to endure painful surgery. Amputees often report pain in their amputated, missing limbs. VR environments can be used to immerse the patient within an environment to better cope with their phantom pains.[158] VR also works as a distraction. In the future, children and adults alike will be given a VR headset from their doctor when receiving an injection to distract them from the impending prick. Hermes Pardini Laboratories and Vaccination Centres piloted the use of a VR headset that immerses the patient into a fictional, gamified environment.[159] The patient is distracted in a virtual world to the point that they are often unaware they have received their injection.

Physical Therapy

VR can track human movement, allowing patient movements to be monitored and analyzed. VR gyms have been opened in San Francisco and Ohio.[160] Rehabilitation will become gamified, for instance, through kicking a virtual ball or catching a ball. Recovery exercises can be delivered and tracked in a VR environment and retold to the patients if they were not to get it right.

Cognitive Rehabilitation

The application of VR and AR in fears and phobias is apparent. Performed as a medical treatment, patients can be gradually exposed, known as graded-exposure therapy. Data from sensors will be assessed to ensure patient safety and develop best practice.

Cognitive function can be improved for patients struggling to perform everyday tasks. For instance, through monitoring the performing of patient tasks in a VR environment, a doctor could determine declining memory loss and identify areas of concern or priority. Similarly, patients with injuries to the brain or those that struggle with tasks can have digital environments created to represent real-life scenarios. Patients can practice tasks and regain or develop cognitive function. Patient engagement can be monitored and analyzed to observe areas of difficulty or reduced attention. There are few randomized controlled trials that have been conducted into VR in cognitive rehabilitation. However, some applications of VR are effective in treating cognitive deficits in people with neurological diagnoses.

Nursing and Delivery of Medicine

VR and AR will develop to become part of standard medical training. Both types of reality can be used to improve outcomes by enriching the information available to the individual. For instance, Alder Leys

Children's Hospital in England uses a VR headset and 360-degree video to train student doctors how to deal with stressful situations. Critical decision points can be reviewed and analyzed with peers and staff. Virtual reality will be used more to learn medical specialties such as performing operations and anatomy. Operations can be practiced with an expert monitoring and giving feedback in real time. The Royal London Hospital conducted the world's first VR surgery in 2016,[161] enabling viewers to watch in 3-D 360. The learning experience is unparalleled and has the potential to disrupt medical training, particularly for countries that have limited healthcare resources. In the future, doctors will use augmented and mixed reality to enrich their environment. Surgeons will have access to critical information in real time through mixed realities, delivered to their vision through glasses. Nurses, for instance, could use augmented or mixed reality to identify veins in the arm for a blood test. The Interventional Cardiology Center at Tufts Medical Center in Boston uses VR to introduce the facility to prospective patients with anxiety before a procedure.[162]

Virtual Appointments and Classrooms

As the cost for virtual and merged reality devices becomes more affordable, virtual appointments will become commonplace. Virtual appointments remove the inconvenience of attending a clinic, saving time, environmental resources and focusing healthcare professional time to where required. Virtual appointments will become as familiar as webinars, enabling stakeholders to be present without traveling. As well as one-to-one appointments, virtual and merged realities provide an immersive and engaging experience suited for learning. Virtual reality will become a staple tool for education and training—for both healthcare professionals and patients alike.

Blockchain

Transitioning of traditional patient health records to the EHR is considered to be enormous progress for healthcare. The digitalization of patient records mitigates some of the traditional risks of centralized data stores. However, this model still places the medical records in the hands of the provider. Blockchain technology, popularized by the bitcoin cryptocurrency, has the potential to revolutionize data access, privacy, and trust. Currently, blockchain is yet to be deployed for mainstream healthcare.

Blockchain, fundamentally a collection of data records, is a piece of software formed by the combination of several preexisting technologies that provide blockchain with its characteristic features: an immutable, distributed public ledger whose authenticity can be verified by anyone: those who validate the data on the ledger are rewarded with value, which helps create trust in a trustless environment; distributed peer-to-peer control, which provides a high level of security; and the ledger can be programmed to trigger automatic transactions in the form of smart contracts, allowing for a widespread application of this technology.

PayPal, Visa, and Mastercard, for instance, act as central authorities for financial transactions: trusted institutions that act as intermediaries. The land registry acts as the trusted store for details on homeownership. These centralized databases are subject to being hacked or manipulated. Blockchain technology aims to solve data management, privacy, and security issues, improving interoperability and easing the flow of data between doctors, hospitals, healthcare systems, and insurance providers through the use of a decentralized, immutable database. Patients are demanding increased access to their medical health records, with iOS 11.3 even including an EHR feature.[163] Data from EHRs IoT, wearables, and devices can be used within the ledger, in a trusted, secure, transparent, and interoperable environment.

The blockchain is an immutable database that is stored and maintained by all those using it. Each new transaction or piece of data is encrypted and then approved by a particular proof-of protocol (consensus, work, stake) by other nodes on the networks that authenticate the transaction or verify the stored piece of data.

Each node on the network has an identical copy of the blockchain; and thus, the transaction is permanently recorded and linked to previous records. The links, known as hashes, are traceable back to the very first block in the blockchain. Therefore, any attempts to tamper with a block in the present would require the transaction and all related blocks to be also altered, on all records, distributed among the nodes holding a copy of the ledger simultaneously. The longest chain of events is considered valid. Blockchain implements crypto-economic and game-theory techniques. Any attempt to create a rival chain would need to be created faster than the current version of the truth to be accepted. This would require tremendous computing, energy, and resource commitments. Miners compete to validate blocks. This good behavior is expensive regarding electricity and computing power, and so miners are incentivized to validate blocks, rewarded with bitcoin (BTC; in the BTC scenario). Through creating a common, distributed, immutable database of healthcare information, doctors and healthcare providers have the potential to access medical data from any system, with improved security and privacy, less administration, and better sharing of results.

- Single points of failure are eliminated by decentralizing and encrypting the data.
- The blockchain is democratized. Anyone can contribute or store a version of the truth.
- Through ensuring consensus on events, the most likely version of the truth is held.

- A transparent and auditable ledger of events is provided through time-stamping.
- Game theory, crypto-economics, and hashing incentivize good behavior and ensure the events are without censorship.

Blockchain technology has many applications in healthcare, with most technology currently in pilot or proof of concept stage.

Verifying the Supply Chain

The first application of blockchain is verifying the supply chain. As the blockchain is created in a chronological and auditable manner, it can act as a means of verification for components of every link in the supply chain. The journey of materials or treatments, for instance, can be logged onto a blockchain ledger, which could be used to identify fraudulent activity, anomalies, error, or a break in the chain through data entry or IoT devices. For instance, the cold-delivery chain of delivering insulin from manufacturer to pharmacy could be confirmed to be untampered and appropriately cooled. This is currently being used in developing countries to combat counterfeit medication. It is also being developed to enhance genomic data protection, addressing the privacy challenges of big genomic data.

Incentivized Wellness

Blockchain technology could be used to incentivize wellness: through the use of a cryptocurrency as a digital token of value. Engaging people with health services or a healthier lifestyle could save the global economy a tremendous amount of money in healthcare costs. Health providers or employers typically see the cost benefit of this in the form of savings or profit, and this is rarely passed to the individual. By utilizing blockchain technology,

tokens could be created and distributed to patients through the blockchain to share the value of the savings and be treated as a tradable currency.

Individuals could earn tokens through behaviors such as going to the gym, reaching their step goal, attending education sessions, engaging in mindfulness, completing a particular sports event, or adhering to medication or digital therapeutics. The ecosystem rewards positive behaviors with an asset, or token of value. The value of the token could be fixed. Extending this concept, positive health behaviors could be extended to the point where patients have health token savings accounts that could be used to transact in hospitals.

Patient Record Access

Patients are demanding access to their health record. This poses a significant challenge as to how to best share sensitive medical data with unknown third parties. For third parties, there is also a challenge to verify the integrity of the data while ensuring privacy for the patient. Presenting patient records on an appropriately permitted blockchain would give cryptographic assurance on data quality without any need for human involvement. Providers or consumers uploading health data would generate transactions. Users provide a signature and timestamp, alongside a private key to access data. By using digital signatures, all records stored on the blockchain can be identified and used to create a comprehensive patient health record. The use of digital signatures and cryptographic encryption ensures data travel securely and are accessible only by those with the relevant public keys. By using blockchain technology, each addition to the EHR can be logged, with an immutable, auditable trail of transactions, while ensuring the most current version of the record is used. Patients would be able to verify all attempts to access or process data. Blockchain's decentralized structure enables any approved stakeholder to join the ecosystem, without the need for data integration or manipulation concerns.

Robots

The use of robotics is changing healthcare, although it will be some time before robotic technology will be widely affordable or implemented.

Indeed, it is also unlikely robots will entirely replace humans in the medical setting. Currently, hospitals and healthcare systems are unable to meet the cost of the technology; and robotic assistants have not, and some would say cannot, replace human contact. Over and above its use in the supply chain, robot technologies have potential in the following areas.

Robot-Assisted Surgery

Robot-assisted surgery enables enhanced vision, improved precision, and dexterity. Currently limited in reach, robot-assisted surgery will become commonplace as barriers to entry such as the cost of hardware and training are reduced. In the future, healthcare professionals may have to learn how to use an instrument as well as how to perform the surgery. Robotic-assisted surgery blurs the lines of liability, which could prove to be a barrier to adoption.

Exoskeletons

Robotic exoskeleton technology focuses on enabling patients to perform tasks through the use of an external, integrated device. Ekso Bionics is an organization that provides a wearable exoskeleton that assists spinal cord injury patients to stand and learn to walk. In the future, exoskeletons will become a standard form of rehabilitation and aid human mobility in more activities, and in more environments.

Inpatient Care

Inpatient care can be enhanced through the use of robots for streamlining tasks. Robots could be used to automate tasks such as collecting mail or delivering blood. Delivery robots will deliver medications or important material autonomously. Robots are already being used to disinfect rooms, where the risk to humans is minimized through using a robotic agent.[164] Soon, ingestible agents and smart pills will be able to monitor a patient's internal reaction to treatment; robotic nurses will become commonplace for automatable tasks, such as taking blood. A robot will take your vitals and draw blood by identifying the correct vein with greater accuracy than a human nurse. Robots will be used to support care for the elderly and those with long-term conditions.

Companions

Virtual assistants are becoming more human-like with developments in natural language processing. As discussed previously, humans are already forming bonds with robots. Robots can be used to provide companionship or treat loneliness in cases of mental health, elderly, and long-term care. As robots evolve in their capabilities, robots will be able to perform increasing healthcare duties such as bathing patients, transporting patients, and so on. Robots will be able to monitor vital signs through IoT and wearables and aid the patient in real time.

Drones

Drones have the potential to alter the way medicine is delivered. Drones can be used in hard-to-reach areas, places of conflict, and remote populations to deliver medication, vaccinations, and diagnostics. Drones will be used to deliver medications from the pharmacy, just as they are being used in shopping. Time-sensitive items such as blood, bodily fluids,

and organs can travel shorter distances to go directly to where they need to be, within campuses or across larger distances. Drones can also be used to locate and identify, particularly in remote locations. One use case for drones is as a flying medication toolbox in cases of emergency. Critical care in the minutes after a stroke, trauma, or heart failure is essential to accelerate recovery and prevent death. An ambulance drone prototype, developed by Tu Delft, combined a heart defibrillator, medication, and two-way radio that could be dispatched to a patient to speed up the response before the first responder.[165]

Current restrictions in the use of drones are reducing their potential. It may be some time before flight restrictions, legislation, and functional issues such as drone battery life and drone load are resolved.

Smart Places

Intelligent homes, hospitals, places, and things promise to change the way we live. Ubiquitous connectivity and an expanding sensor base provide a wealth of opportunity to people, patients, and providers alike. Improvements in patient lifestyle and healthcare have contributed to growing numbers of centenarians over the past 30 years. In the United Kingdom alone, the number of people living to 100 increased by 65%.[166] These breakthroughs have come at a cost to healthcare; and smart places, and smart things, enable scalable opportunities to improve health and social care—reducing costs per capita while personalizing the experience to the patient. Rather than going through the process of connecting and integrating apps, sensors, or devices, connected places will use automated sensors that do not require the user's constant attention, collecting tremendous amounts of real-time data. Facial recognition, voice recognition, responsive notifications, and data-based suggestions will become the norm. Aggregation and analytics platforms will streamline and simplify decision-making processes in a rapidly changing environment.

These devices can notify patients and healthcare professionals in real time: forget your Big Brother, this is Big Doctor.

Smart Homes

There will be many ways that homes will become smart in the pursuit of improved patient and population health, which are best illustrated by example. On waking, a wearable sleep monitor will assess the quality of your sleep while you perform mindfulness through your connected watch. Your sleep monitor will even tell you when to go to bed for optimum recovery. After brushing your teeth with a toothbrush that assesses whether you are hydrated or not, you drink your morning coffee, which controls the dosage of insulin released for absorption in your bloodstream. Your fridge will inform you of out-of-date food and ensure none of your potential allergens are found in the food you purchase. Smartphone apps will augment the calories, and nutritional values of foods; and 3-D printers will enable the printing of safe-to-dispense, at-home pharmaceuticals. Virtual assistants and even mirrors will assess for signs of anxiety and depression and assess health biomarkers in natural language. Should you develop a cough, your virtual assistant will tell you that you're coughing more than usual; and with the change in body temperature detected from your smartwatch, your assistant calls the doctor on your behalf and arranges an appointment. As a closed environment, the smart home allows for AI to accommodate for the highly individualized needs of the user. The most important element aspect of smart anything is the nuance when scaling: to understand the context and prioritize local activity.

The future is not as far as it seems. Bolzano in Italy is already working with IBM and various partners to empower aging people to age safely at home.[167] Safety and security are the main priorities, with the project using sensors installed into homes to monitor environmental factors such as temperature, carbon monoxide, and water leaks. Data is pushed to an off-site

control room where, depending on the individual need, family members, volunteers, emergency response staff, and social services are notified.

Smart Hospitals

A smart hospital is one that is connected, much like the smart home. The objective of a smart hospital is to provide clinical excellence, an efficient supply chain, and superb patient experience—facilitated with technology. Smart hospitals will use a continuous learning ecosystem, ranging across many areas including electronic data collection and health records, digital technology, robotics, 3-D printing, unstructured data, and robust analytics. Somewhat ironically, mass use of technology and AI will enable patients to become more actively involved in their treatment decisions. Non-emergency consultations will take place on the Internet, with AI prioritization enabling the right doctor, with the right skills and training, to treat the patient. Treatment will be blended between offline, physical care and digital care—where adherence and accountability can be quantified and maintained. Growing data and large-scale analytics will continue to personalize digital treatments for patients.

On attending a hospital, an automated and streamlined admission reduces patient waiting time. On admission, patients will be tagged with a clinical-grade wearable to track vital signs through the inpatient stay. Metrics are all sent wirelessly to a dashboard visible to your medical team. Any anomalies or causes for concern are detected and prioritized. Hospitals and surgeries will become hubs of data, with hospitals working with business to extract value from the data—regarding improving efficiency, minimizing mistakes, and for improving treatment and device decisions. Healthcare will be delivered as a service, with patients incentivized with cryptocurrency to maintain sensible and healthy lifestyle choices. All relevant data will be anonymized and accessible to digital health partners and internal departments to enable continuous learning from each patient's experience.

Reductionism

The basic premise of reductionism is that by breaking down or simplifying complex biological or medical phenomena into their many parts, one is much more likely to understand a single cause and devise a cure. Reasoning is reductionist. For instance, a mouse can be divided into a skeleton, a circulatory system, a nervous system, a digestive system, and so on. It fits a model—a pattern, theory, equation, formula, or some form of common structure. AI systems also follow the same, reductionist approach; reducing the problem to a defined discipline. For AI to truly develop artificial understanding, the key will be to address humans from a perspective of holism.

As the world continues to immerse itself in digital technology, healthcare AI models are becoming quicker and more accurate, with patient-based care and lifestyle medicine placing the focus back on holistic health. The effect of lifestyle factors, sleep, movement, nutrition, environment, and genetics on health is well known. And we know that the world and people, mind and language, can be unpredictable. As AI develops, problems of irreducible complexity will require solving where the reductionist approach cannot work. The focus is on value, on the patient; treating patients as people and focusing on the patient's wider health, relationships, and goals. As AI grows, so too does the requirement to not lose the human touch.

Innovation vs. Deliberation

While AI and digital health hold significant promise, healthcare is facing significant challenges in coping with the rapid drive to integrate these systems into medicine—where the stakes are high. Meaningful changes in healthcare will come from the right blend of innovation and deliberation. As the boundaries of discovery are tested, the potential of providing more

patients personalized and precise care is realized. At the same time, there is an obligation to explore and anticipate the social, economic, political, and ethical consequences of innovation.

AI's adoption in healthcare is best managed as a gradual development, which would be useful for stakeholders to have time to map and understand the possible consequences of its implementation. Robust standards are required for auditing and understanding AI models to uphold quality assurance and trust. Academic research, clinical evidence, and policymaking are also required on the use of AI systems in healthcare. The development of AI standards requires a diverse and public forum that is academically robust and reviewed periodically in line with the latest evidence base.

Awareness within the medical professional and patient communities is required. Healthcare professionals need basic knowledge of how AI applications function in their respective medical settings to understand the advantages, disadvantages, and consequences of how AI can facilitate and assist them in their everyday tasks. Patients must become accustomed to engaging with AI systems and explore the benefits for themselves. Demonstratable examples of how AI in healthcare is optimizing efficiency, optimizing resources, and improving population health outcomes will strengthen the public perception of AI.

Organizations working with AI need to communicate and engage the general public about the potential advantages and risks of using AI in medicine. Organizations developing AI must work with academia to take the necessary steps to be able to measure the success and the effectiveness of AI systems independently. It is vitally important to encourage affordable AI solutions to democratize AI as the stethoscope of the twenty-first century; and published, validated evidence and literature is essential to trust and adoption. Here's to the most hopeful era for healthcare yet.

CHAPTER 8

Case Studies

AI is improving the healthcare experience, bringing success to those who can leverage and adapt to a new health and care delivery paradigm. The proceeding case studies provide unique and engaging perspectives of the use of big data, AI, and machine learning within healthcare. Real-life descriptions of organizational approaches to data-identified healthcare problems demonstrates the instant value within available data.

Several enthusiastic and responsive healthcare organizations submitted a case study for inclusion within the book. The focus of case studies was real-world demonstrations of AI, machine learning, data, and IoT usage that were novel and innovative in their provision. As well as describing what success looked like, and how it was achieved, the aim was to bring light to the inertias and developed learnings encapsulated within systems.

Case Study Selection

A call for case studies was placed on Twitter and we received over 60 case study applicants. Case studies were only accepted on the proviso that the application was deployed in some guise, whether pilot or otherwise, and case studies were supported with evidence to demonstrate value from a

financial and/or patient perspective. The criteria used to evaluate case studies was the following:

- Focus on delivering precision medicine
- The case study focused on a machine learning or data-driven model that was in existence, in production, and deployed
- Case studies focused on learnings from the projects
- Case studies required authors to share adequate information on the project scope and conclude on the organizational, business, clinical, or financial impact

The case studies chosen for inclusion have been validated and contain applications that exist in healthcare today. Case studies comprise the following:

- AI for Imaging of Diabetic Foot Concerns and Prioritization of Referral for Improvements in Morbidity and Mortality—Harkrishan Singh, Gro Health
- Outcomes of a Digitally Delivered, Low-Carbohydrate, Type 2 Diabetes Self-Management Program: 1-Year Results of a Single-Arm Longitudinal Study—Laura Saslow, Jim Aikens, and David Unwin
- Delivering a Scalable and Engaging Digital Therapy for Epilepsy—Charlotte Summers, Diabetes Digital Media
- Improving Learning Outcomes for Junior Doctors Through the Novel Use of Augmented and Virtual Reality—Alder Ley Children's Hospital
- Big Data, Big Impact, Big Ethics: Diagnosing Patient Risk from Data—Dom Otero, Diabetes.co.uk

Conclusion

Case studies contained within this chapter showcase the potential of data-driven systems, AI, and machine learning in healthcare. The case studies selected showcase a tremendous vision, focus, and drive to improve patient engagement, outcomes, and clinical success. Showcased case studies serve as stimulating examples on the topics of the following:

- The use of data to drive decision-making and provide patient- and population-scale precision medicine
- Advantages and pitfalls of developing machine learning models
- Data collection, analysis, and governance
- Developing intelligent agents and the blend between digital and offline healthcare
- Ethics of AI and morality
- The future of healthcare and digital technology

Case Study: AI for Imaging of Diabetic Foot Concerns and Prioritization of Referral for Improvements in Morbidity and Mortality

Harkrishan Singh, Head of Machine Learning, Gro Health, United Kingdom

Background

Over 400 million people globally have diabetes, a condition of high blood sugar (glucose) call hyperglycemia. The two main types of diabetes are type 1 and type 2. Chronic hyperglycemia can cause recognized complications of diabetes including foot problems. Due to prolonged

high blood glucose levels, people with diabetes may have slower wound healing times. They may also have reduced nerve functioning called peripheral diabetic neuropathy. This means the nerves that usually carry pain sensation to the brain from the feet do not function as well, and it is possible for damage to occur to the foot without feeling it.

Treading on something, wearing tight shoes, cuts, blisters, and bruises can all develop into diabetes foot ulcers. Narrowed arteries can also be a consequence of diabetes, and reduced blood flow to the feet can also impair the foot's ability to heal properly. When the foot cannot heal, a foot ulcer can develop. Foot ulcers are a common complication of poorly controlled diabetes, forming as a result of skin tissue breaking down and exposing the layers underneath. Chronic poor blood supply can contribute to the risk of requiring amputation of toes, feet, and lower legs. Approximately 2.5% of the population with diabetes have foot ulcers at any given time, which in England alone is 68,000 people. £1 in every £140 of NHS spending is on ulceration and amputation. Those who wait longer for specialist care tend to have more severe ulcers, and their ulcers are less likely to have healed at 12 weeks. Only 20% of patients are referred to a foot protection service within the National Institute of Clinical Excellence (NICE)-recommended time frame of 1 day.

In 2014–2015, 1 in every 390 people with diabetes were admitted to the hospital for an amputation; and 1 in 33 people with diabetes were hospitalized due to a foot ulceration. The combined NHS cost for inpatient care was £322 million (86% on ulcer-related admissions).

Inpatient improvement programs delivered by video to promote foot checks and reduce ulcers have been successful. Ipswich Hospital NHS Trust saved an estimated 19 foot ulcers and 571 diabetic bed days per year, saving £246,000 annually. The cost benefit from improved foot self-care, referral, and diagnosis is substantial. NICE guidance recommends those with an active foot problem should be referred to a multidisciplinary team (MDT) or foot protection service within one working day and triaged within one further working day. The National Diabetes Foot Audit (NDFA) found that 20% of 9,137 patients followed were seen in this time frame.

Cognitive Vision

Cognitive vision or image recognition is a machine learning technique designed to mimic how the human brain functions. With this technique, machine learning models are taught how to recognize visual components that make up an image. Through learning on large datasets of images and identifying patterns, models are able to make sense of their input and determine relevant tags and classifications. Image recognition is not an easy task. To successfully perform image recognition, neural networks are the technique of choice. However, even still, the computation resource cost is expensive in practice. For instance, a 20-pixel by 20-pixel image received by a neural network would receive over 400 inputs. Although this sounds achievable in a typical hardware setup, images of greater size, say 1,000-pixels by 1,000-pixels, would require a powerful computational resource to process the higher number of parameters and inputs. In a traditional neural network, each pixel would be linked to a single neuron. This is computationally expensive.

Proximity of pixels within images has a strong relationship with their similarity. Convolutional neural networks specifically make use of this feature. Rather than treating two nearby pixels as distinct, convolutional neural networks assume there is more likely a relationship between these pixels than two that are further apart. Through bypassing less significant connections between pixels, convolution solves the computational and time problems faced by traditional neural networks. Convolutional neural networks improve the computational resource required for image recognition through filtering relationships by proximity.

Clinical evidence suggests that diagnostic tests and risk stratification can predict the risk of ulceration and amputation, with early referral reducing amputation rates and time to heal. Solutions for diabetic foot care must be broader than the standard healthcare worker dependent service and be deliverable at a lower cost for the patient and healthcare system. Together with my colleague Yang Wang, we applied state-of-the-art AI

techniques to automatically detect signs of diabetic foot concerns, alert the patient, and support communication and assessment by the healthcare team. Image recognition technology for diabetic foot concerns would be revolutionary, particularly considering this machine learning experience could improve ulceration healing time, cost, and risk of amputations—as this is related to how early ulceration is detected. The model has achieved high accuracy in detecting subtle areas of bruising and cracked heels, which can soon develop into areas of foot ulceration.

A cloud-based system was developed to allow remote use of the application such that any device with an Internet browser was able to upload an image alongside a standardized protocol. The output would be a list of any concerns detected, the referral pathway most suited to the concern, and education, if relevant. The model is constantly refined through a machine learning feedback loop that optimizes and retrains the model periodically with random subsets of labelled data, including the data received from the web-based system.

Project Aims

The aims of our project are to further develop and validate this machine learning technology, including the deep learning neural network used to predict the risk of foot concerns and provide precise medicine to patients. The convolutional neural networks used for multiple predictions can become more precise through its use and feedback loop from users. As the user base is largely healthcare professionals and carers, the learning can be scaled up and validated by the healthcare professional for appropriate triage. The project also aims to deliver a service provision to patients, which is only possible when the model has reached maximum, and ideally near-absolute, accuracy. Reducing error is a key priority for the project. Grading schemes and standardized protocols support the model to develop greater precision and recall. The ability to identify ulceration and other foot concerns earlier would accelerate diagnosis, treatment,

and save the lives of many. A cloud-based system would enable easier sharing of the application and its outcomes. A catchall event log was developed to identify all events taking place within the ecosystem. The platform includes behavior health mentoring and integrated tracking for blood glucose, blood pressure, mood, food, weight, and sleep through Bluetooth-enabled devices, wearables, and self-inputted data.

AI models provide boundless opportunity in predictive analytics. By analyzing and generalizing on thousands of images, medical images are being turned into science. Precision medicine is the objective of evidence-based medicine. There are few ways as engaging and immersive as digital, which enables healthcare to collect enormous amounts of data; analyze and provide useful feedback to inform patients; and prevent, predict, diagnose, and treat disease. By the very nature of imagery, de-identification for aggregated learning is technically simple. However, the requirement to be able to explain the mechanisms of deep neural network prediction is technically difficult. A more general understanding of the application alleviates most stakeholder concerns, which can be explained through demonstration. As a novel innovation, there has been a requirement to explain how models are generated, validated, and continually evaluated. Regulatory approval has proved more laborious than previous projects due to the loose definitions that exist around product categorization. Error minimization has been a priority, and regulatory approval has focused on ensuring patient safety. The risk of false positives was determined to be more likely: early ulceration detection means quicker healing and less risk of amputation. Thresholds of concerns were developed within the model, with any doubt referred to a human.

Outputs of the project are being documented through clinical papers to highlight the value of the project to patients, healthcare providers, and global healthcare providers. Perhaps of most significance, the outcomes will demonstrate how people with diabetes can achieve significant health improvements by preventing and detecting diabetic foot problems.

Challenges

There were a number of inertia points that were observed through the duration of the project:

- **Quality of data**

Any model can only be as good as the data that is used to create it. Data completeness and correctness are paramount for a robust machine learning algorithm. There were several significant challenges with the datasets collected that facilitated the development of robust data governance processes. First, lack of labelling meant that data was unlabelled and often unclear on receipt. Second, handwritten notes that supported image scans were often illegible. Incorrect spelling or incomplete doctor notes complicated the data validation process. The cleaning and vetting data advanced a dataset that was established as a source of truth for the model.

- **Data governance**

The requirement for national and international regulatory approval required appropriate, transparent data governance processes to be in place. Governance-covered areas included data architecture, archiving, management, metadata management, privacy, security, and validation. Standards, best practices, and procedures were developed for data acquisition, verification, and validation. A central repository held all relevant data and metadata, with access restricted by credentials to ensure appropriate user management.

A governance trail for all decisions was noted as per the requirements for regulation, which provided an opportunity for reflective learning as the project ended. Stringent data governance ensured the integrity of data used within the project.

- **Performance vs. “explainability”**

Neural networks are notorious for being difficult to explain, particularly as models become more complex. A key requirement of this project was transparency in algorithmic decisions. As the model developed, iterations of learning and backpropagation tailored nodes within the network that became unexplainable. Model performance and accuracy was secondary to model transparency.

- **Ethical governance**

Strict data ethics and information ethics governance procedures were developed alongside data governance processes to ensure ethical and moral use of data. The machine learning model was used by healthcare professionals to predict and diagnose risk of foot ulceration. Processes around the effect of false positives and true negatives in particular were noted, as misdiagnosis was established to be the greatest ethical concern. The effects of ulceration and healing times can be greatly reduced through swift diagnosis and action. As the application was used to support the actions of healthcare professionals, users were informed as to the limitations of the model, and it was explained how model accuracy was based on validation datasets.

- **Closing the loop**

The developed model predicted the likelihood of foot ulceration and other diabetic foot concerns. As this was to be used in the stages of prevention and diagnosis, a feature that was requested was the ability to validate and confirm model predictions. To do so, the healthcare professional using the application would either confirm diagnoses or refer the patient to a specialist for diagnosis. The process of feeding back patient outcomes into the system enabled the system to learn in near real time. Data feedback from patients once referred to the hospital was difficult and cumbersome to gather.

- **Stakeholder understanding**

An unforeseen challenge was the lack of stakeholder understanding and sporadic resistance to technology-driven healthcare solutions. Key stakeholders were typically biased by previous experiences, negative press, or threatened by concerns of an age of robots taking jobs from humans. This was overcome with staff stakeholder training days; third-party experts and internal stakeholders explained how AI and machine learning could be used to optimize the tasks conducted by healthcare professionals, improve the resource burden, and mitigate concerns of a near-future robot race. The project scope was developed with the input of a diverse, multifaceted stakeholder group to ensure all relevant parties were vested in the project.

- **Adoption strategy**

As a project that started through R&D funding, there was little incorporation of the pilot project into the wider organizational direction. Stakeholder comfortability and adoption of the novel innovation was progressed through workshops and hands-on experience. Workshops focused on mitigating stakeholder concerns and on the ease, efficacy, and resource-light environment of the predictive system.

Conclusions

Cloud-based intelligence that is sensitive and accommodating to differences in photography hardware and formats can be combined with patient clinical, demographic, behavioral, and genomic data to develop a precision-medicine framework for detection of diabetic foot concerns. The prototype was welcomed by healthcare professionals, with several remarking on elements that improved their own knowledge on the topic. Healthcare professionals felt more confident in diagnosing foot concerns when partnered with the digital tool. Digital assistance can confirm and support decisions and challenge healthcare professionals on how they arrive at their decisions. Applied across the spectrum of medical images and data collected by healthcare and research organizations, the technology will affordably be able to enable large-scale prognostic, risk stratification, and best treatment decisions.

From our experience, having more data is almost always more important than the algorithm itself. Overcoming the concerns of text digitalization and label recognition alongside medical imagery proves to be a challenging area within this field. Handwriting and human error have been the biggest problems to mitigate to reduce false recall. Conducting a large-scale clinical trial is the next phase of the project. A pilot project assessing the feasibility of the project is pending funding approval.

Case Study: Outcomes of a Digitally Delivered, Low-Carbohydrate, Type 2 Diabetes Self-Management Program: 1-Year Results of a Single-Arm Longitudinal Study

LR Saslow, PhD, University of Michigan, Ann Arbor; JE Aikens, PhD, University of Michigan, Ann Arbor; DJ Unwin, FRCGP, Principal in General Practice, The Norwood Surgery, Southport, UK

[The following details the Low Carb Program digital health intervention from Diabetes Digital Media. The intervention is used by over 300,000 people in 190 countries.]

Background

Type 2 diabetes has serious health consequences including blindness, amputation, stroke, and dementia; and its annual global costs are more than \$800 billion. Although typically considered a progressive, nonreversible disease, some researchers and clinicians now argue that type 2 diabetes may be effectively treated with a carbohydrate-reduced diet.

Although typically considered a progressive, nonreversible disease, some researchers and clinicians now argue that type 2 diabetes may be effectively treated with a carbohydrate-reduced diet, which could improve type 2 diabetes management and potentially even lead to remission.[168] Indeed, previous research with carbohydrate-reduced diets for type 2 diabetes do show improved outcomes (such as glycemic control, weight loss, and reductions in the use of hypoglycemic medications), for both very low-carbohydrate diets (roughly 20% or fewer of total dietary calories derived from carbohydrates)[169, 170, 171] or lower carbohydrate diets (roughly 40% or fewer of total dietary calories derived from carbohydrates). [172, 173] Although dietary interventions have historically been in person,

online programs can be just as effective for some participants, as suggested by research that has examined diet and lifestyle interventions in adults with prediabetes.[174] Therefore, it is perhaps not surprising that the beneficial results of carbohydrate-reduced diets for people with type 2 diabetes (glycemic control, weight loss, and reductions in the use of hypoglycemic medications) have been replicated using online programs. [175, 176]

Objectives

Our objective was to evaluate the 1-year outcomes of a digitally delivered Low-Carb Program (LCP), a nutritionally focused, 10-session educational intervention for glycemic control and weight loss for adults with type 2 diabetes. The program reinforces carbohydrate restriction using behavioral techniques including goal setting, peer support, and behavioral self-monitoring.

Methods

The study used a quasi-experimental research design comprised of an open-label, single-arm, pre- and post-intervention using a sample of convenience. From adults with type 2 diabetes who had joined the program and had a complete baseline dataset, we randomly selected participants to be followed for 1 year (N = 1,000; mean age 56.1, SD 15.7, years; 59% [593/1,000] women; mean HbA1c 7.8, SD 2.1, %; mean body weight 89.6, SD 23.1, kg; taking an average of 1.2 diabetes medications).

The LCP is a completely automated, structured, 10-week health intervention for adults with type 2 diabetes. Participants are given access to nutrition-focused modules, with a new module available each week over the course of 10 weeks. The modules are designed to help participants gradually reduce their total carbohydrate intake to < 130 g

per day to meet their self-selected goals. The program encourages participants to make behavior changes based on “Action Points” or behavior-change goals at the end of each module. These goals are supported with resources that are available to download including information sheets, recipes, and suggested food substitution ideas. The LCP online platform also includes digital tools for submitting self-monitoring and device-driven data on a number of different variables including blood glucose levels, blood pressure, mood, sleep, food intake, and body weight. Weekly automated feedback is provided to users based on their use of the program through e-mail notifications, and participants are notified when the next week’s module has been opened. Lessons are taught through videos, written content, or podcasts of varying lengths. The program stresses the importance of regular contact with the participants’ healthcare providers for adjustments in medications in weeks 1, 2, and 10. After the 10 weeks of modules have been opened, participants continue to have access to the education content as well as the ability to continue to track their health (glycemic control, weight). The content and strategies used in the program build off of prior research and theory. For example, evidence suggests that goal setting can act as an effective behavior change strategy used to improve adherence to lifestyle intervention programs in obesity management programs. The program therefore encourages participants to select a goal at the beginning of the program (such as to lose weight, reduce medication dependency, or make healthier choices for their whole family). Participants are also prompted to consider how their health would benefit from attaining their goal. Throughout the program, participants are periodically prompted to consider how close they are to attaining their goal. The program further reinforces behavior change through integrated tracking, whereby program users are encouraged to track their health data including mood, food intake, blood glucose levels, weight, sleep, and HbA1c. According to the control theory of behavior change, monitoring goal progress—that is, evaluating one’s ongoing

performance relative to the standard—and responding accordingly is critical to goal attainment. Recent findings suggest that program interventions that elevate the frequency of progress monitoring are likely to induce behavior change.

Results

Of the 1,000 study participants, 708 (70.8%) individuals reported outcomes at 12 months, 672 (67.2%) completed at least 40% of the lessons, and 528 (52.8%) completed all lessons of the program. Of the 743 participants with a starting HbA1c at or above the type 2 diabetes threshold of 6.5%, 195 (26.2%) reduced their HbA1c to below the threshold while taking no glucose-lowering medications or just metformin. Of the participants who were taking at least one hypoglycemic medication at baseline, 40.4% (289/714) reduced one or more of these medications. Almost half (46.4%, 464/1,000) of all participants lost at least 5% of their body weight. Overall, glycemic control and weight loss improved, especially for participants who completed all 10 modules of the program. For example, participants with elevated baseline HbA1c ($\geq 7.5\%$) who engaged with all 10 weekly modules reduced their HbA1c from 9.2% to 7.1% ($P < .001$) and lost an average of 6.9% of their body weight ($P < .001$).

Observations

- The engagement platform used by patients for this study was only available on the Web and not as a mobile app. It would be expected that a mobile app would improve engagement.
- The criteria for inclusion within the project was the requirement to be over 18 and the ability to speak English.

- The percentage of individuals with an HbA1c level of < 6.5% increased from 26% (257/1,000) to 50% (503/1,000). This degree of control, when achieved through pharmacotherapy, is often accompanied by weight gain and risk for hypoglycemic events.[177] As the now famous Action to Control Cardiovascular Risk in Diabetes (ACCORD) study reported, intensive hypoglycemic medical therapy “increased mortality and did not significantly reduce major cardiovascular events.”[178]
- A limitation was the rate of delivering the entire intervention, as only 528 (52.8%) completed all modules. However, a high rate (70.8%) reported 12-month outcomes. This could be due to the program being launched in November, with Christmas and seasonal activities affecting the rate of completion. On the other hand, given that this program was entirely automated and had a wide reach, a large number of individuals were able to complete the program.

Conclusions

Especially for participants who fully engage, an online program that teaches a carbohydrate-reduced diet to adults with type 2 diabetes can be effective for glycemic control, weight loss, and reducing hypoglycemic medications.

Case Study: Delivering A Scalable and Engaging Digital Therapy for Epilepsy

Charlotte Summers, Chief Operations Officer, Diabetes Digital Media, United Kingdom.

Background

Epilepsy is a neurological condition of brain. Different epilepsies are due to many different underlying causes. The causes can be complex, and sometimes hard to identify. A person might start having seizures because they have one or more of the following: a genetic tendency, structural changes in the brain, or genetic conditions. Epilepsy is sometimes referred to as a long-term condition, as people often live with it for many years, or for life. Although generally epilepsy cannot be “cured,” for most people, seizures can be “controlled” (stopped) so that epilepsy has little or no impact on their lives. Treatment is often about managing seizures in the long term. Most people with epilepsy take antiepileptic drugs (AEDs) to stop their seizures from happening. However, there can be side effects with such medications, and there are other treatment options for people whose seizures are not controlled by AEDs, including the ketogenic diet.

Implementing the Evidence Base

The positive effects and therapeutic mechanisms of dietary ketosis and that of a ketogenic diet (KD) on human physiology have been well documented in literature. The KD is a high-fat, low carbohydrate (usually less than 50 g/day), adequate protein diet (Paoli et al. 2013), whose metabolic effects originates back in the 1960s; however the therapeutic effects of KDs can be traced back to the early 1920s when it

was successfully used in the treatment of epilepsy (Kessler et al. 2011). Ever since, the potential clinical utility of KDs has been investigated by several studies, resulting in an accumulation of scientific evidence on the therapeutic role of KDs on various physiological disorders. Several studies have shown that the KD does reduce or prevent seizures in many children whose seizures could not be controlled by medications. Over half of children who go on the diet have at least a 50% reduction in the number of their seizures. Some children, usually 10–15%, even become seizure free. The Epilepsy Society support the KD as a treatment option for patients over 12 months old. Research highlights the therapeutic effects of the KD on patients with epilepsy as primarily reduced medication and reduced seizures. In a recent Cochrane systematic review of the evidence regarding the effects of KDs, Levy and Cooper found no randomized controlled trials (RCTs).

This demonstrates that epilepsy can be managed; and in some cases, patients can live seizure-free lives through the sustained application of a ketogenic way of eating. Through reducing sugar in the diet and at the same time improving blood glucose control, people with epilepsy can achieve significant health benefits such as reducing their risk of seizures and number of medications.

Sensor-Driven Digital Program

The Ketogenic Program for Epilepsy (KPE) is a structured education behavior change program for people with epilepsy. The KPE empowers users to sustainably adopt a lower carbohydrate lifestyle with the appropriate personalized education, health-tracking facilities, support, health mentoring, and resources to maintain a safe and sustainable lifestyle. As a result, people with epilepsy could expect to improve glycemic control and as a result reduce the incidence of seizures, reduce medication dependency, and improve confidence in managing their condition.

Health data is collected in real time through blood glucose monitoring, medication monitoring, and seizure tracking modules within the application. Unstructured data is particularly useful in understanding sentiment and the patient's psychology. Research activity will take place over the course of 12 months, reporting on 3-month, 6-month, 9-month, and then 1-year epidemiological health outcomes and engagement data followed by an 18-month and 2-year follow-up to demonstrate efficacy and adherence to the program over the short to medium term. This will be the world's largest and longest study on patients engaged with a digital platform epilepsy. It is possible that with the right funding this could be conducted as a randomized controlled diet, which would be the world's first in KDs delivered for epilepsy through a digital platform.

Research

As part of this project, we seek to learn the most efficient and effective methods to implement and accelerate at-scale delivery of our technology to the UK's NHS population and international bill payers. The project is led by an award-winning and experienced team, recognized for innovating digital health, who are determined to revolutionize the health and well-being of people across the globe. The technological challenge is to create a scalable, engaging, and effective solution for global implementation. This challenge is mitigated through the project extending on the infrastructure of DDM's (Diabetes Digital Media) Low Carb Program for type 2 diabetes, with clinically validated outcomes and 71% retention at 1 year. Surprisingly, out of 300,000 patients globally registered within the Low Carb Program for type 2 diabetes, there are 3,112 people with epilepsy (0.5% of the UK population)—showing the readiness for an online, nutrition-focused intervention for this population.

DDM intends to conduct the world's first RCT on patients following a ketogenic protocol with the Ketogenic Program pioneering in the space of

digital health innovation. The potential cost-saving impact to the NHS is significant. Patients who are empowered and have improved health and well-being are also more active members of their communities and have lower social care needs.

Project Impact

The greatest impact will be felt by patients with epilepsy and their families through empowering patients to manage their epilepsy. Evidence demonstrates that a KD in people with epilepsy can be maintained by infants and adults with improvements in number of seizures and medication consumed. Evidence also demonstrates positive effects on cardiovascular disease risk. The health improvements and potential cost savings are sizeable. In addition to this, the project is of tremendous value to the UK economy and the NHS, with reduced complications, improved mental and emotional health, reduced hospital admissions, and reduced medication having a direct impact on the budget of the UK Treasury.

Positive social impacts experienced by the patient have a repercussion for wider communities, businesses, and the UK economy. As the result of improved patient and population health, organizations will benefit from fewer sick days and reduced absenteeism, improved mental health and improved perceived quality of life, as well as a reduction in the perceived burden of managing epilepsy. The positive impact is also received by the clinical healthcare community—through reducing physician and healthcare professional burden and enhancing the evidence base.

Preliminary Analysis

The objective of the research arm of this project is to evaluate the 1-year outcomes of a digitally delivered Ketogenic Program: a nutritionally focused, 16-session educational intervention for epilepsy control through provision of a KD. The program assists in patients placing themselves in

ketosis using behavioral techniques including goal setting, peer support, and behavioral self-monitoring. Interesting correlations between elevated blood glucose and mood enabled predictive algorithms to be developed to identify symptoms of impending seizures, alerting patients through notifications to take action. Especially for participants who fully engage, an online program that teaches a carbohydrate-reduced diet to adults with epilepsy can be effective for glycemic control, weight loss, and reducing medication and the number of patient seizures. An RCT is required to understand the clinical impact of such a digital therapeutic.

Case Study: Improving Learning Outcomes For Junior Doctors Through the Novel Use of Augmented and Virtual Reality

Paul Duval, University of Liverpool; Vidhi Taylor-Jones, University of Liverpool

Background

Junior doctors arrive at a university eager to help mankind, yet often they are unaware of how to deal with high-pressure, critical scenarios. Dealing with medical concerns appropriately—whether emergencies or not—requires years of experience and learning. Over time, exposure in doctor surgeries, hospitals, and clinics develops the self-assurance, readiness, and experience required for junior doctors to realize their vocation. High-pressure scenarios, such as a patient having a cardiac arrest, require pinpoint precision accuracy in decision-making and actions. For some time after first arrival to a university, junior doctors are not ready to make such decisions. This is where augmented and virtual reality provides an opportunity to engage students in learning that readies them for life as a doctor. What's

more, with 1,500 students in any one year, group trips to hospitals are expensive and logistically difficult to manage.

The project was recognized for its pioneering nature in 2016, when it won an Association for the Study of Medical Education (ASME) Education Innovation Award.

Aims

The aims of the project were manyfold:

- Increase the exposure to simulation and training environments for junior doctors without the use of simulation centers.
- Give inexperienced doctors the experience they require in a noncritical environment.
- Demonstrate effective and ineffective clinical practice.
- Encourage students to consider their own practices and thought processes.
- Immerse students in a clinical learning environment and identify whether the use of virtual reality enabled them to become familiar with the healthcare setting, patient unpredictability, and provided the confidence to cope with high-pressure scenarios.

Project Description

It was agreed that the project have as wide a reach as possible. This eliminated virtual reality hardware such as the HoloLens or Oculus and focused the delivery onto mobile devices. The Samsung Gear and Google Pixel were able to handle the delivery of the content with the use of a headset to extend the view into virtual reality. This was logical considering

almost every student owns a phone, whereas the use of HoloLens and Oculus is still relatively niche.

Simulations of a variety of clinical experiences were filmed with 360-degree videography using both doctors and patients as actors. These simulations were short 5- to 10-minute videos that featured decision moments where students needed to make an active choice—including a patient suffering from cardiac arrest and a patient with schizophrenia becoming hysterical. These videos were integrated into a virtual reality application that enabled students to explore a clinical setting in virtual reality, engage in the environment, and encouraged them to make a variety of decisions during the process.

Conclusions

- Virtual reality simulation provided a prime and safe clinical environment for junior doctors to explore their practice.
- In particular, the technology allowed students to focus on mistakes—unifying approaches to medicine across the cohort.
- The technology enabled students to learn anywhere, anytime: at home, school, or library.
- Virtual reality simulations were accessible by all medical students with the appropriate hardware.
- Students were able to become more active in decision-making in lectures.
- The virtual reality experience encouraged discussion and learning.

- It was cheaper and logistically easier to manage than student trips to hospital wards.
- The use of this innovative technology prepares junior doctors for the demands of the vocation.
- Research and exploration into the appropriate pedagogies to support virtual reality learning is required as the tools become more freely available.

Case Study: Big Data, Big Impact, Big Ethics: Diagnosing Disease Risk from Patient Data

Dominic Otero Head of Engineering, Diabetes.co.uk, United Kingdom

Background

The variety and velocity of patient data has exponentially increased over recent decades. This data has not progressed the robustness of decision-making. This leads to questions as to why decision-making is not becoming more data-driven and what can be done to validate and expedite its adoption.

Diabetes.co.uk is the world's largest diabetes community, based at the University of Warwick, and the platform welcomes over 45 million visitors a year. As the world's diabetes community, the organization provides a range of services that focus on the patient-data ecosystem to empower patients to manage and control their health. By collecting patient data, or real-world evidence, over time, the ecosystem is created such that it offers insights, education, and feedback on particular aspects of the patient's health.

Platform Services

The platform collects a wealth of data. From the moment of platform access, behavioral and engagement metrics are collected, with patients progressing onto registering for the platform, opting in to part with over 120 variables that cover their diabetes and overall health status. Once inside the platform, a variety of data types are collected—unstructured data in the form of conversations, interactions, engagements, and behavior; and structured data in the form of self-reported, device-driven, wearable, IoT, and clinical health record data. The platform provides a number of services to patients and bill payers.

Medication Adherence, Efficacy and Burden

Medication adherence, adverse effects, and wastage are concerns that pharmacology seeks to overcome. Patients with diabetes are typically on at least one medication on entry to the platform. Within the platform, patients evaluate their medications at regular intervals, sharing their own opinions on the use of medicine in the real-world, and report more structured data such as side effects, adherence, efficacy, and burden of using medications. This data is mapped geographically, by condition, medication, and other markers and fed back in a de-identified, aggregate manner to pharmacological and academic partners for research purposes. Data is used by partners to identify new drugs and understand interactions, usage, adherence, and real-world adverse effect prevalence. The same data is also used to improve patient safety. For instance, the platform alerts members on particular medications or using particular devices should there be an issue or recall.

Community Forum

The Diabetes.co.uk community forum is the world's largest support community for people with diabetes and provides an unrivalled insight into the global diabetes population. The platform welcomes people with all types of diabetes, carers, and healthcare professionals. The concept is simple: real people with diabetes talking about their own experiences and speaking to others who are in a similar position. Patient engagement is high; there is over 2 million years of cumulative experience among members, who spend over 17 minutes each visit on the platform. As one of the world's most novel digital communities, data and respective insights from the online forum and dissemination of academic findings are constantly shared with our patient-support network, medical advisory board, partners, and other interested networks.

A first phase study by Royal Holloway, University of London, in 2017 investigated the role of online health communities in patient self-management, assessing the following:

- How patients appraise knowledge from online health communities in managing diabetes; and
- How a patient's healthcare professionals perceive and can potentially use such knowledge to innovate their clinical practice.

The study concluded the following:

1. The Diabetes.co.uk forum is a catalyst of innovation.
2. The Diabetes.co.uk forum empowers patients, meaning
 - i. A good relationship with their healthcare professional
 - ii. High confidence or self-efficacy in managing their condition
 - iii. Feeling less emotionally burdened

This was supported with data from 13,000 patients who highlighted the patient benefit through platform usage:

- 53% improved blood glucose control
- 58% improved quality of life
- 59% improved confidence in managing diabetes
- 65% improved dietary choices
- 76% improved understanding of diabetes

AI prioritization of patient interactions

To meet regulatory requirements, any health platform must report patient health or safety concerns immediately on discovery. This would prove cumbersome for a developed health platform such as Diabetes.co.uk if completely human-led for several reasons:

- Monitoring of medical discussion groups and interactions requires human-intensive staffing to spot concerns for regulation purposes.
- Lots of time is wasted if humans were to do all the checking.

It was decided that to meet our obligations, the use of AI bots would speed up the process, enabling humans to spend time with the concerns that matter. This was realized through the development of a neural network-based classifier to determine sentiment and concern based on contents, user profiles, and frequency of other key metrics. The classification system went through several hundred rounds of validation, taking a subset of the huge data lake that was available from the millions of posts. The classifier was able to spot 99.8% of all concerns, with humans

required to intervene in 0.2% of cases. There were a number of benefits to the AI prioritization:

- It saved a considerable amount of human time.
- Human time was spent on the problems that were of most concern.
- The flagging of discussions occurred in near real time rather than the 15–20 minutes it would take a human to do so.

This information also assisted the customer support teams to efficiently deal with user queries by knowing exactly what a conversation was about, typical responses to such a query, and what to say to a user to ensure a favorable outcome for all parties involved. Additionally, by gauging the sentiment of users and storing it, user data can be implemented to build additional systems in the future.

Real-World Evidence

Through the analysis of patient data, we can determine a discrepancy between official and real-world health statistics. For instance, 14% of people with diabetes in the United Kingdom are expected to have a mental or emotional health concern, whereby the platform sees 44% of users with this health condition. This highlights that further exploration and understanding of the patient population is severely needed to ensure critical aspects of patient care are not neglected. Interestingly, partners have become more interested in real-world evidence as the years have progressed, partially to do with concerns over research funding and conflicts of interest.

Ethical Implications of Predictive Analytics

Predictive analytics enables intelligent classifiers to predict the risk of disease. The collection of a varied patient dataset—including conversations, health markers, demographics, and behavior—enables the development of a sophisticated model that can be used to refine predictions further. One of the unintended consequences of the project was discovered when deploying an AI model that predicted the risk of a patient having pancreatic cancer. Pancreatic cancer typically effects 1 in 10,000 people and is frequently misdiagnosed. Often, patients receive a diagnosis when it's too late. A pancreatic cancer AI model was developed using a clinical dataset of health metrics, engagement, and behavior data and first predicted the likelihood of poor patient metabolic health. Should patients meet the threshold for poor metabolic health, the model evaluates the risk of the patient having pancreatic cancer.

Communicating this information to users was a first-time problem for the organization. All relevant stakeholders, including all members of the Medical Advisory Panel, were prompted for their input on this topic. Predictive analytics had previously been used to inform patients of possible hypertension or increased risk of conditions based on lifestyle. However, informing patients of their risk of pancreatic cancer was considered to be unethical, particularly as we did not want to create additional mental health concerns or anxieties.

How should a user be told of a possible life-changing concern? The key action for this communication is to nudge the patient to see their doctor or physician as soon as possible. The factor for success in this case was determined to be user friction. The less friction there is between the user and communicating party, the more successful the engagement. Stakeholders from behavioral economics were vital in testing and confirming an engaging communication to inform the user that their data fell outside expected norms without raising fear. To do this, the prediction was framed against patient and population demographic and medical

data. Rather than being told they were the only patient to see anomalous data, patients were reassured that others have also seen similar patterns and been successful in improving their health.

Integration of the IoT

Sensors in wearables such as accelerometers, altimeters, and heart rate and skin temperature monitors provide a tremendous amount of useful information. For instance, motion, sweat molecules, and conversations are all useful and relevant data points. Such metrics can be used to take a snapshot of metabolic health, cardiovascular health, and mental health. For this to be conducted at scale, there must be little to no friction between the patient and providing parties. The less friction there is between the user and whom they choose to share data with, the more successful the product. Wearables are becoming increasingly used in incentivized wellness propositions. Validating patient behaviors and general physiology facilitates the personalization of premiums and outcomes-based reimbursements.

Patients have the option to connect their devices and applications to the Diabetes.co.uk architecture. What is interesting is that younger users typically adhere to wearable usage and find it easy, while elderly users typically did not. This is of relevance to insurance and incentivized wellness propositions, as most claims come from older policyholders.

Conclusions

The realization of precision medicine is not without technical and ethical challenges. IoT and conversational and wearable data are just examples of a variety of data that can be collected and used to empower patients in real time, facilitating improvements in health and well-being. However, with data and the application of AI, come ethical concerns, particularly in the prediction of future disease risk. The lack of processes and governance on such topics is urgently required as the abilities of AI mature.

Technical Glossary

- Accuracy

A metric used in machine learning that is used to evaluate the performance of a machine learning model.

- Backpropagation

Used by neural networks, the updating of biases and weights based on error. This occurs when the estimated output exceeds a defined error threshold. The error at output is propagated back into the network to update values of neurons.

- Bagging

A machine learning technique of creating multiple models on subsets of the same data and combining them to improve overall prediction.

- Bootstrapping

The technique of splitting data into multiple subsets with replacement. Each sample is known as a bootstrap sample.

- Big data

A large corpus of data—which is typically both structured and unstructured.

TECHNICAL GLOSSARY

- Columnar

The storing of data in columns rather than rows (as per many traditional database structures). Columnar data storage improves speed due to reduced disk load and improved data transfer.

- Continuous variable

A variable that can take any infinite value between its minimum and maximum.

- Cost function

A cost function is a performance metric that is used to measure the error of a machine learning model.

- Dataframe

A dataframe is a labelled data structure that can contain columns of different types of data.

- Dataset

A dataset is a collection of examples.

- Data transformation

The process of converting data from one form or type to another

- Discrete variable

A variable that can take a finite number of values.

- Example

A collection of features.

- Exploratory Data Analysis (EDA)

EDA is a data science technique that seeks to understand data insights through statistical analysis and graphical visualization.

- False negative

A model output that is incorrectly predicted as false.

- False positive

A model output that is incorrectly predicted as true.

- Falcon

A management framework for Apache, used for governance of data in Hadoop clusters.

- Feature

A feature is a data attribute and its value. As an example, skin color is brown is a feature where skin color is the attribute and brown is the value.

- Feature selection

The process of choosing the features required to explain the outputs of a statistical model while excluding irrelevant features.

- Flume

A distributed service that collects, aggregates, and transfers large, real-time data into the HDFS (Hadoop Distributed File System).

- F-score

An evaluation technique that combines precision and recall into a measure of classification effectiveness.

- Goodness of fit

This term refers to how well a model fits an observation set and summarizes any differences between observed and expected values.

- Hadoop

Hadoop is a commonly used, open source framework developed by Apache that caters for distributed processing of big datasets.

- Hadoop Common

Common libraries, modules, and extensions that support Hadoop modules.

- Hadoop MapReduce

Hadoop MapReduce is a design framework for software development that facilitates the processing of large datasets in parallel.

- HBase

Refers to a Hadoop database that enables the storage and management of sparse data.

- Hive

Hive is an open source library that enables Hadoop tasks to be programmed using SQL. This provides a relational database storage structure.

- JSON

JSON, or JavaScript Object Notation, is a lightweight data file type that facilitates quicker retrieval of data, particularly in web-focused applications, as it is easier to parse and generate.

- R

R is a software and programming language for statistical computing and graphics. Project R has been designed as a data mining tool, while R programming language is a high-level statistical language that is used for analysis.

- Impala

An open source, massively parallel processing database for Apache that enables data querying from HDFS or HBase.

- Matlab

Matlab is the language many university students begin with. It is useful for fast prototyping, as it contains a large machine learning repository.

- Matrix

A matrix is an array used for representation of data.

- MongoDB

MongoDB is a NoSQL database, which are used to store unstructured data with no particular schema.

- Model selection

The process of choosing a statistical model for machine learning from known models.

- NewSQL

NewSQL databases maintain the integrity of typical relational SQL databases while providing the scalable performance of NoSQL structures.

- NoSQL

A modern approach to databases that is useful for managing data that changes frequently, or data that is unstructured or semi-structured in nature. Rows can all have their own set of unique column values. NoSQL has been driven by the requirements to provide better performance in storing big data. The architecture has better write performance, takes less storage space with compression, and reduces operational overhead.

- Overfitting

Overfitting occurs when a machine learning model performs well on a training dataset but does not perform as well on a validation set. The model essentially learns the training data.

- Parameters

Attribute values that control the output and behavior of a machine learning system.

- *P* value

On the basis the null hypothesis is true, the *P* value refers to the probability of receiving a value equal to or greater than the observed value.

- Ranger

A data management and monitoring framework to ensure data security.

- Regression

This is a supervised learning technique in which the output is a real value.

- Sqoop

An extremely useful, open source tool to facilitate the transfer of data between Hadoop and traditional relational database systems.

- Spark

Spark is a simple programming model that can be used with Java, Scala, Python, and R that enables large-scale data processing applications to be written quickly.

- SQL

SQL (structured query language) is a language used for managing data held in a traditional database management system.

- Tensor

A tensor is an object similar to a vector that is represented as an array that can hold data in N dimensions. A tensor is a generalization of a matrix in N -dimensional space.

- TensorFlow

An Alphabet-backed, open source library of data computations optimized for machine learning that enables multilayered neural networks and quick training.

- True negative

A model output that is correctly predicted as false.

- True positive

A model output that is correctly predicted as true.

TECHNICAL GLOSSARY

- Underfitting

Underfitting occurs when a machine learning model is not able to capture the signal of the data and will have poor performance on both training and validation data.

- XML

An acronym for eXtensible Markup Language, XML is another form of flat data file designed to make the importing, exporting, generation, and movement of data easier.

References

1. <https://www.forbes.com/sites/bernardmarr/2017/09/22/12-ai-quotes-everyone-should-read/>
2. <http://www-01.ibm.com/software/data/bigdata/>
3. <https://www.forbes.com/sites/parmyolson/2018/06/28/ai-doctors-exam-babylon-health/#7076587912c0>
4. <https://diabetestimes.co.uk/nhs-green-light-for-low-carb-programme/>
5. Ibid
6. Lessons learned review of the WannaCry Ransomware Cyber Attack, NHS <https://www.england.nhs.uk/wp-content/uploads/2018/02/lessons-learned-review-wannacry-ransomware-cyber-attack-cio-review.pdf>
7. The One-Way Mirror: Public attitudes to commercial access to health data, March 2016
8. Outdated NHS technology hindering cost savings, retrieved <https://www.clarityworkforcetech.com/blog/outdated-nhs-technology-hindering-cost-savings/>
9. Sharing of Data: a Patient Perspective: data on file, Diabetes.co.uk, 2018
10. The Digital Edge, Gartner, <http://www.gartner.com/imagesrv/books/digital-edge/TheDigitalEdge.pdf>
11. Big Data: present and future, Mircea Răducu TRIFU, http://www.dbjournal.ro/archive/15/15_4.pdf

REFERENCES

12. A Study on the Role of Big Data in Social Media Tadepalli, Sarada Kiranmayee, March 2015
13. Identifying big data dimensions and structure, 2017 4th International Conference on Signal Processing, Computing and Control (ISPPCC)
14. Big Data, Analytics Market To Hit \$203 Billion In 2020, Information Week
15. <http://ns1758.ca/winch/winchest.html>
16. <http://www.diabetes.co.uk/facesofdiabetes>
17. Twitter data accurately tracked Haiti cholera outbreak, Nature, 2012
18. IBM Insights on Business, <https://www.ibm.com/blogs/insights-on-business/consumer-products/2-5-quintillion-bytes-of-data-created-every-day-how-does-cpg-retail-manage-it/>
19. Detecting influenza epidemics using search engine query data, Nature Vol 457, 19 February 2009
20. <https://www.mckinsey.com/industries/healthcare-systems-and-services/our-insights/the-big-data-revolution-in-us-health-care>
21. Open graph, <http://ogp.me/>
22. Risk stratification of early admission to the intensive care unit of patients with no major criteria of severe community-acquired pneumonia: development of an international prediction rule, published online 2009 Apr 9. doi: 10.1186/cc7781
23. Allocating scarce resources in real time to reduce heart failure readmissions: a prospective, controlled study
24. <https://www.optumlabs.com/>
25. Overview, eHealth : Digital health and care, EU, retrieved 14/1/2017

26. What health systems can learn from Kaiser Permanente,McKinsey & Company. Retrieved 14/1/2017
27. <https://diabetestimes.co.uk/nhs-green-light-for-low-carb-programme/>
28. <https://www.blueshieldca.com/bsca/about-blue-shield/media-center/calindex-launch-080514.sp>
29. <http://www.modoo-med.com/>
30. http://institute.swissre.com/research/library/health_monitoring_dave_wang.html
31. The Influence of Wireless Self-Monitoring Program on the Relationship Between Patient Activation and Health Behaviors, Medication Adherence, and Blood Pressure Levels in Hypertensive Patients: A Substudy of a Randomized Controlled Trial, published online 2016 Jun 22. doi: 10.2196/jmir.5429
32. <https://www.wired.com/story/science-says-fitness-trackers-dont-work-wear-one-anyway/>
33. Your Fitness Tracker Isn't the Best Way to Measure Heart Rate, Cleveland clinic, <https://health.clevelandclinic.org/your-fitness-tracker-isnt-the-best-way-to-measure-heart-rate/>
34. US AID, Use of technology in the Ebola Outbreak in West Africa, https://pdf.usaid.gov/pdf_docs/PA00K99H.pdf
35. Improved Response to Disasters and Outbreaks by Tracking Population Movements with Mobile Phone Network Data: A Post-Earthquake Geospatial Study in Haiti, published: August 30, 2011, <https://doi.org/10.1371/journal.pmed.1001083>
36. TAKMI Bringing Order to Unstructured Data, <http://www-03.ibm.com/ibm/history/ibm100/us/en/icons/takmi/>

REFERENCES

37. Evidence Generation Guide for Apps and Wearables Developers: Study Designs Including Applied Examples, YHEC, NHS, retrieved 14/5/2018
38. LowCarbProgram.com
39. Brighton and Sussex University Hospitals NHS Trust, Data Protection and Audit Report, ICO, 2018
40. Cyber security: Attack of the health hackers, FT, retrieved 18 June 2018
41. Hackers can make your pacemaker or your insulin pump kill you – and the NHS needs to respond to that threat Cyber-security is not just about keeping data safe - it's about protecting patients' health; Ara Darzi, <https://www.independent.co.uk/voices/hackers-medicine-nhs-cyber-attack-medical-device-pacemaker-wifi-a8032251.html>
42. Half of web traffic is now encrypted, TechCrunch, retrieved Feb 28, 2018
43. Arthur Lee Samuel, IEEE Computer Society, <http://history.computer.org/pioneers/samuel.html>
44. Artificial intelligence expedites breast cancer risk prediction, Houston Methodist, retrieved February 29, 2018
45. Ibid
46. Artificial Intelligence: A Modern Approach (International Edition)
47. <http://www.contrib.andrew.cmu.edu/~mndarwis/ML.html>
48. AlphaGo, Deepmind, <https://deepmind.com/research/alphago/>
49. <http://www.independent.co.uk/life-style/gadgets-and-tech/news/facebook-artificial-intelligence-ai-chatbot-new-language-research-openai-google-a7869706.html>
50. Data Mining: Opportunities and Challenges by John Wang

51. <https://link.springer.com/article/10.1007/BF00993309>
52. Stochastic Gradient Boosting, Jerome H Friedman, 1999
53. World's best Go player flummoxed by Google's 'godlike' AlphaGo AI, The Guardian, <https://www.theguardian.com/technology/2017/may/23/alphago-google-ai-beats-ke-jie-china-go>
54. Porter Stemmer, <http://people.scs.carleton.ca/~armyunis/projects/KAPI/porter.pdf>
55. An Introduction to Genetic Algorithms, Jenna Carr
56. Diabetes prevalence, <https://www.diabetes.co.uk/diabetes-prevalence.html>
57. Risk Factors for Incidence of Cardiovascular Diseases and All-Cause Mortality in a Middle Eastern Population over a Decade Follow-up: Tehran Lipid and Glucose Study. Published online 2016 Dec 8. doi: 10.1371/journal.pone.0167623
58. Machine Learning and Data Mining Methods in Diabetes Research. Published online 2017 Jan 8. doi: 10.1016/j.csbj.2016.12.005
59. Ibid
60. Population-Level Prediction of Type 2 Diabetes From Claims Data and Analysis of Risk Factors. Published Online:1 Dec 2015. <https://doi.org/10.1089/big.2015.0020>
61. Development and validation of risk assessment models for diabetes-related complications based on the DCCT/EDIC data. <https://doi.org/10.1016/j.jdiacomp.2015.03.001>
62. An interpretable rule-based diagnostic classification of diabetic nephropathy among type 2 diabetes patients. Published online 2015 Jan 21. doi: 10.1186/1471-2105-16-S1-S5

REFERENCES

63. Using a multi-staged strategy based on machine learning and mathematical modeling to predict genotype-phenotype risk patterns in diabetic kidney disease: a prospective case-control cohort analysis. Published online 2013 Jul 23. doi: 10.1186/1471-2369-14-162
64. Risk Stratification of Patients with Diabetes: Depression and Distress. Published online on June 22, 2018.
65. Random Search for Hyper-Parameter Optimization, James Bergstra, Yoshua Bengio; 13(Feb):281–305, 2012.
66. In Brief: The P Value: What Is It and What Does It Tell You? Published online 2010 May 25. doi: 10.1007/s11999-010-1402-9
67. <https://www.pkvitality.com/>
68. <https://www.diabetes.co.uk/artificial-pancreas.html>
69. Bioprinting of skin constructs for wound healing. Published online 2018 Jan 23. doi: 10.1186/s41038-017-0104-x.
70. The Role of Online Health Communities in Patient Empowerment An Empirical Study of Knowledge Creation and Sharing on Diabetes. co.uk, Dr Roberta Bernardi, Royal Holloway, https://pure.royalholloway.ac.uk/portal/files/29019759/Findings_report.pdf
71. The Dataverse Project: An Open-Source Data Repository and Data Sharing Community, Research Data Access & Preservation Summit, Atlanta, GA
72. Burden: mortality, morbidity and risk factors, WHO, http://www.who.int/nmh/publications/ncd_report_chapter1.pdf
73. Global health workforce shortage to reach 12.9 million in coming decades, WHO
74. The Debate on the Moral Responsibilities of Online Service Providers. Sci Eng Ethics. 2016 Dec;22(6):1575-1603. Epub 2015 Nov 27.

75. Health Tech and Mobile Apps: an Analytical Report.
76. The Information of Intelligence, Scientific American.
77. Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach, <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>
78. Apple health data used in murder trial, BBC. <https://www.bbc.co.uk/news/technology-42663297>
79. Did senators questioning Facebook's Mark Zuckerberg understand the internet? The Guardian, <https://www.theguardian.com/technology/video/2018/apr/12/did-senators-questioning-facebooks-mark-zuckerberg-understand-the-internet-video>
80. EU GDPR <https://www.eugdpr.org/>
81. Ibid
82. DDM Health Interventions, <https://www.diabetes.co.uk/healthcare-professionals/index.html>
83. Cambridge Analytica kept Facebook data models through US election, <https://www.theguardian.com/uk-news/2018/may/06/cambridge-analytica-kept-facebook-data-models-through-us-election>
84. Data on file, Diabetes.co.uk, January 2018.
85. WeLiveSecurity reveals 73,000 unprotected security cameras with default passwords, WeLiveSecurity, 2016.
86. Google Nest hacked 'in 15 seconds' as reality bites for Internet of Things, <https://www.scmagazineuk.com/google-nest-hacked-in-15-seconds-reality-bites-internet-things/article/1480564> retrieved 8 August 2017.

REFERENCES

87. The 'three black teenagers' search shows it is society, not Google, that is racist, <https://www.theguardian.com/commentisfree/2016/jun/10/three-black-teenagers-google-racist-tweet>
88. Microsoft deletes 'teen girl' AI after it became a Hitler-loving sex robot within 24 hours, <https://www.telegraph.co.uk/technology/2016/03/24/microsofts-teen-girl-ai-turns-into-a-hitler-loving-sex-robot-wit/>
89. The accuracy, fairness, and limits of predicting recidivism. Science Advances 17 Jan 2018: Vol. 4, no. 1, eaao5580, DOI:10.1126/sciadv.aao5580
90. Four Kinds of Ethical Robots, Philosophy Now.
91. CES 2018: LG robot Cloi repeatedly fails on stage at its unveil, <https://www.bbc.co.uk/news/technology-42614281>
92. Breast screening error 'shortened up to 270 lives' – Hunt, <https://www.bbc.co.uk/news/health-43973652>
93. How Uber Deceives the Authorities Worldwide, NY Times, <https://www.nytimes.com/2017/03/03/technology/uber-greyball-program-evade-authorities.html>
94. How They Did It: An Analysis of Emission Defeat Devices in Modern Automobile, <http://cseweb.ucsd.edu/~klevchen/diesel-sp17.pdf>
95. <https://lis.epfl.ch/>
96. I, Robot (1950) Asimov.
97. South Korean Robot Ethics Charter 2012.
98. The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, 2017.
99. Mastering the Game of Go without Human Knowledge –DeepMind.

100. Google's 'deceitful' AI assistant to identify itself as a robot during calls, <https://www.theguardian.com/technology/2018/may/11/google-duplex-ai-identify-itself-as-robot-during-calls>
101. How to Hold Algorithms Accountable, MIT Technology Review, 2016.
102. Turing Test Success Marks Milestone, University of Reading. <http://www.reading.ac.uk/news-and-events/releases/PR583836.aspx>
103. Engineer 'Marries' Robot He Built And It's Totally Not Creepy At All https://www.huffingtonpost.co.uk/entry/zheng-jiajiarobot-marriage_us_58e3c701e4b0d0b7e1651098
104. Trump digital director says Facebook helped win the White House, <https://www.theguardian.com/technology/2017/oct/08/trump-digital-director-brad-parscale-facebook-advertising>
105. It's not a drug, but it may as well be": Expert opinions on whether kids are addicted to tech, CABA, Understanding Addiction. <https://www.caba.org.uk/helpandguides/information/understanding-drug-addiction>
106. Action on Addiction, <https://www.actiononaddiction.org.uk/>
107. Man Dies After 3 Day Internet Binge in Taipei, <https://www.smh.com.au/technology/man-dies-afterthreeday-internet-gaming-binge-20150117-12sg0a.html>
108. McDonald's Introduces Screen Ordering and Table Service, NY Times, <https://www.nytimes.com/2016/11/18/business/mcdonalds-introduces-screen-ordering-and-table-service.html>
109. Data on file, Diabetes.co.uk 2018.
110. Gray Matters: Too Much Screen Time Damages the Brain, Victoria L. Dunckley M.D.

REFERENCES

111. Too Much Screen Time Associations, https://www.gla.ac.uk/news/headline_586422_en.html
112. AI Predicts Lifespan, <https://www.adelaide.edu.au/news/news92622.html>
113. The AI arms race: China and US compete to dominate big data, FT, <https://www.ft.com/content/e33a6994-447e-11e8-93cf-67ac3a6482fd>
114. UK can lead the way on ethical AI, says Lords Committee, Parliament.UK.
115. Meet Sophia, the world's first robot to get citizenship, <https://www.scmp.com/business/global-economy/article/2123889/meet-sophia-worlds-first-robot-get-citizenship>
116. Individual NHS doctors receiving £100,000 per year from drugs firms, <https://www.telegraph.co.uk/news/2016/06/30/individual-nhs-doctors-receiving-100000-per-year-from-drugs-firm/>
117. Tseng, Eric K., and Lisa K. Hicks. "Value based care and patient-centered care: divergent or complementary?" *Current hematologic malignancy reports* 11.4 (2016): 303-310.
118. Buurtzorg at the Curse of Plenty Conference at SwissRe, Zurich, May 2018. http://institute.swissre.com/events/14th_CRO_Assembly_Risk_Focus.html#tab_3
119. Cancer is a Preventable Disease that Requires Major Lifestyle Changes. DOI: 10.1007/s11095-008-9661-9.
120. The Effects of Stress on Physical Activity and Exercise. *Sports Med.* 2014 Jan; 44(1): 81–121. doi: 10.1007/s40279-013-0090-5.

121. Saslow LR, Summers C, Aikens JE, Unwin DJ Outcomes of a Digitally Delivered Low-Carbohydrate Type 2 Diabetes Self-Management Program:1-Year Results of a Single-Arm Longitudinal Study, JMIR Diabetes. DOI:10.2196/diabetes.9333.
122. Evidence based medicine: what it is and what it isn't BMJ 1996; 312 doi: <https://doi.org/10.1136/bmj.312.7023.71>
123. Bad Pharma: How Medicine is Broken, and How We Can Fix It, Ben Goldacre.
124. Data on file, Diabetes.co.uk MyHealth, 2018.
125. Genetic Factors Are Not the Major Causes of Chronic Diseases, PLoS One 2016; 11(4): e0154387.
126. Diabetes in South Asians, Diabetes.co.uk <https://www.diabetes.co.uk/south-asian/>
127. McGrath, Scott. (2018). The Influence of 'Omics' in Shaping Precision Medicine. EMJ Innovations. 2. 50-55.
128. Researchers Develop World's Smallest Wearable Device, <http://www.mccormick.northwestern.edu/news/articles/2018/01/researchers-develop-worlds-smallest-wearable-device.html>
129. Researchers develop smart tattoos for health monitoring, <https://phys.org/news/2017-09-smart-tattoos-health.html>
130. Scientists create painless patch of insulin-producing beta cells to control diabetes, <https://phys.org/news/2016-03-scientists-painless-patch-insulin-producing-beta.html>
131. Data on file, Diabetes.co.uk
132. World Aging Populations, United Nations.
133. Chronic Disease: A Vital Investment, WHO.
134. Chronic Disease Prevention and Health Promotion, CDC

REFERENCES

135. <https://activeprotective.com/>
136. <https://www.nominet.uk/emerging-technology/harnessingpower-iot/digital-future/>
137. Medication Adherence: WHO Cares? - Mayo Clin Proc. 2011 Apr;86(4): 304–314. DOI: 10.4065/mcp.2010.0575
138. Ibid.
139. Ingestible sensors <https://www.proteus.com/>
140. Apple Watch can detect diabetes with an 85% accuracy, TechCrunch, <https://techcrunch.com/2018/02/07/the-applewatch-can-detect-diabetes-with-an-85-accuracy-cardiogramstudy-says/>
141. 325,000 mobile health apps available in 2017, Research 2 Guidance, <https://research2guidance.com/325000-mobile-health-appsavailable-in-2017/>
142. Connected Asthma, Asthma UK <https://www.asthma.org.uk/globalassets/get-involved/external-affairs-campaigns/publications/connected-asthma/connected-asthma---aug-2016.pdf>
143. GP at hand: NHS Doctor Appointments Online, <https://www.gpathand.nhs.uk/>
144. National Diabetes Audit, 2016-2017
145. No health without mental health: A cross Government mental health outcomes strategy for people of all age, Department of Health, UK.
146. Mental Health, WEC, <https://toplink.weforum.org/knowledge/insight/a1Gb0000000pTDbEAM/>
147. Facebook Messenger hits 100,000 bots, <https://venturebeat.com/2017/04/18/facebook-messenger-hits-100000-bots/>
148. <https://www.ncbi.nlm.nih.gov/pubmed>

149. The role of artificial intelligence in precision medicine, <https://doi.org/10.1080/23808993.2017.1380516>
150. <https://cardiogr.am/>
151. The Cost Of Developing Drugs Is Insane. <https://www.forbes.com/sites/matthewherper/2017/10/16/the-cost-of-developingdrugs-is-insane-a-paper-that-argued-otherwise-wasinsanely-bad/#4797ab762d45>
152. The drug-maker's guide to the galaxy, Asher Mullard, Nature.com.
153. Banning cluster munitions, WHO
154. 3-D bioprinter to print human skin, Universidad Carlos III de Madrid – Oficina de Información Científica
155. Alder Hey Children's Hospital, <https://alderhey.nhs.uk/>
156. Genetic disorders <http://www.geneticdisordersuk.org/aboutgeneticdisorders>
157. Virtual reality and pain management: current trends and future directions, Pain Manag. 2011 Mar; 1(2): 147–157. DOI: 10.2217/pmt.10.15
158. UoP News, Five Ways VR is improving Healthcare, University of Portsmouth.
159. Lobo, <http://lobo.cx/vaccine/>
160. Icoros, <https://www.icaros.com/>
161. Cutting-edge theatre: world's first virtual reality operation goes live, <https://www.theguardian.com/technology/2016/apr/14/cutting-edge-theatre-worlds-first-virtual-realityoperation-goes-live>
162. <https://www.tuftsmedicalcenter.org/patient-care-services/>

REFERENCES

163. <https://www.apple.com/healthcare/health-records/>
164. Service robotics: do you know your new companion? Framing an interdisciplinary technology assessment, Published online 2011 Nov 3. doi: 10.1007/s10202-011-0098-6
165. Ambulance Drone, Tu Delft <https://www.tudelft.nl/en/ide/research/research-labs/applied-labs/ambulance-drone/>
166. Record number of centenarians in UK, BBC News, <https://www.bbc.co.uk/news/education-37505339>
167. Case Study: IoT for the OAP: new technology for an ageing population, IBM.com
168. Feinman RD, Pogozelski WK, Astrup A, Bernstein RK, Fine EJ, Westman EC, et al. Dietary carbohydrate restriction as the first approach in diabetes management: critical review and evidence base. *Nutrition* 2015 Jan;31(1):1-13.
169. Hussain TA, Mathew TC, Dashti AA, Asfar S, Al-Zaid N, Dashti HM.
170. Effect of lowcalorie versus low-carbohydrate ketogenic diet in type 2 diabetes. *Nutrition* 2012 Oct;28(10):1016-1021.
171. Nielsen JV, Joensson EA. Low-carbohydrate diet in type 2 diabetes: stable improvement of bodyweight and glycemic control during 44 months follow-up. *Nutr Metab (Lond)* 2008 May 22;5(1):14.
172. Yancy W, Foy M, Chalecki A, Vernon M, Westman E. A lowcarbohydrate, ketogenic diet to treat type 2 diabetes. *Nutr Metab (Lond)* 2005 Dec 01;2(1):34.
173. Garg A, Bantle JP, Henry RR, Coulston AM, Griver KA, Raatz SK, et al. Effects of varying carbohydrate content of diet in patients with non-insulindependent diabetes mellitus. *JAMA* 1994 May 11;271(18):1421-1428.

174. Reaven GM. Effect of dietary carbohydrate on the metabolism of patients with non-insulin dependent diabetes mellitus. *Nutr Rev* 1986 Feb;44(2):65-73.
175. Bian R, Piatt G, Sen A, Plegue MA, De MM, Hafez D, et al. The effect of technology-mediated diabetes prevention interventions on weight: a metaanalysis. *J Med Internet Res* 2017 Mar 27;19(3):e76.
176. Saslow L, Mason A, Kim S, Goldman V, Ploutz-Snyder R, Bayandorian H, et al. An online intervention comparing a very low-carbohydrate ketogenic diet and lifestyle recommendations versus a plate method diet in overweight individuals with type 2 diabetes: a randomized controlled trial. *J Med Internet Res* 2017 Feb 13;19(2):e36.
177. McKenzie A, Hallberg S, Creighton B, Volk B, Link T, Abner M, et al. novel intervention including individualized nutritional recommendations reduces hemoglobin a1c level, medication use, and weight in type 2 diabetes. *JMIR Diabetes* 2017 Mar 07;2(1):e5.
178. Heller S. Weight gain during insulin therapy in patients with type 2 diabetes mellitus. *Diabetes Res Clin Pract* 2004 Sep;65 Suppl 1:S23-S27. Action to Control Cardiovascular Risk in Diabetes Study Group, Gerstein HC, Miller ME, Byington RP, Goff DC, Bigger JT, et al. Effects of intensive glucose lowering in type 2 diabetes. *N Engl J Med* 2008 Jun 12;358(24):2545-2559.

Index

A

Abduction reasoning, 61

Action to Control Cardiovascular
Risk in Diabetes
(ACCORD), 320

Additions, 242

Agent, 77–78

Agent learning, 232

Aggregate data, 216, 220

Algorithm accuracy evaluation,
108–109

AlphaGo, 235

Anonymized data, 216

Application programming interface
(API), 33

Apriori algorithm, 161–162

Artificial intelligence (AI), 281

application, 3

chatbot, 13

components, 76

concepts, 4

conversational AI, 282–283

definition, 1, 4

development, 8

healthcare potential, 15

bias, 17

data-driven strategies, 15

data governance, 16

fragmented data, 15

security, 16

limited memory systems, 6

machine learning

approach, 4

making better doctors (*see*

Doctors and healthcare
professionals, AI)

mining records, 281

natural language conversation, 2

reactive systems, 6

self-aware AI, 7

data-led approach, 8

image recognition, 8

medical health sector, 7

theory of mind, 7

turing test, 2

*Artificial Intelligence: A Modern
Approach* (AIMA), 78

Artificial neural networks

(ANNs), 149, 151–152

Artificial pancreas, 208

Asimov's Laws, 233–234

Assistance Publique-Hôpitaux de
Paris (AP-HP), 44

Association for the Study of Medical
Education (ASME), 326

INDEX

Association rule learning
 method, 160
Atrial fibrillation (AFIB), 38
Auditing, 253–254
Augmented reality (AR), 290–291
Autonomy, 78

B

Behaviors, 242
Bias, 110, 228
Bias correction, 227–228
Big data
 challenges
 data growth, 62
 fragmentation, 65
 infrastructure, 62
 policies and governance, 65
 quality, 63
 quality expertise, 63
 resistance, 64
 security, 63–64
 sources, 63
 strategy, 65
 timeliness of analysis, 66
 visualization, 66
 characteristics, 26
 convert, information, 53
 descriptive
 analytics, 54–55
 diagnostic analytics, 55
 predictive analytics, 55–57
 prescriptive analytics, 58
 definition, 26

 evolution and analytics
 fraudulent transactions, 51
 integration, 52
IoT
 evolution, 32
 patient behavior, 33
 sources, 33
project, 71–72
tools, 72
 Hadoop, 72
 HDFS, 73
 MongoDB, 73
 NoSQL databases, 73
validity, 41
value, 37
 AFIB, 38
 flu symptoms, 38
variability, 41
variety, 31
 cyber-physical devices, 32
 legacy data, 34
 structured data, 32
velocity, 34
 clinical value, 35
 cloud computing, 34
 data sources, 35–36
 traditional server, 34
 wearable biomedical
 sensors, 37
veracity
 factors, 39
 human-based problem, 41
 six Cs, 39–40
visualization, 42

- volume, 28
 - datafication, 30
 - fitness trackers, 29
 - Moore's Law, 29
 - storage factors, 31
- 3 Vs, 27
- 10 Vs, 28
- Big datasets, 221
- Bioprinting of skin constructs, 208
- Blockchain technology, 294
 - hashes, 295
 - incentivized wellness, 296–297
 - medical data, 295–296
 - patient record access, 297
 - verifying supply chain, 296
- Blood glucose device and mobile app, 215
- Blood glucose levels, 207

C

- C45algorithm, 134
- Candidate hypothesis, 106
- Car insurance, 219
- Catchall approach, 30
- Catchall event log, 311
- Chronic disease, prevalence, 209
- Classification and Regression Trees (CART), 134
- Cloud-based intelligence, 315
- Clustering, 158
- Code of ethics, 249, 251
- Cognitive rehabilitation, 292
- COMPAS algorithm, 227

- Connected medicine, 269
 - accessible diagnostic tests, 277
 - digital health and therapeutics, 278–279
 - disease and condition management, 274–275
 - education, 279
 - incentivized wellness, 280
 - medication adherence, 277
 - remote monitoring, 276
 - smart implantables, 278
 - virtual assistants, 275–276
 - wearables and embedded sensors, 270–271
- Conversational AI, 282–283

D

- Data analysis, 223
- Data and information governance
 - asset, 66
 - availability, 68
 - content, 69
 - MDM, 69
 - quality, 68
 - security, 68
 - stewardship, 67
 - use cases, 69
 - benefits, 69
- Data bias, 226
- Data consent, 213
- Data controller, 217
- Data ethics, 66, 210
- Datafication, 30

INDEX

- Data governance policy, 246
- Data management, 39
- Data mining, 92–93
- Data ownership, 215, 220
- Data preparation, machine
 - learning
 - aggregation, 100
 - attribute selection, 99
 - dataset size, 99
 - decomposition, 100
 - transforming data, 100
 - weightings, 101
- Data processor, 217
- Data science, 9–10
- Data science ethics, 210
- Dataset, 82
- Data sharing, 220–221
- Decision-making, 207
- Decision tree
 - benefits, 131
 - ID3, 133
 - test dataset, 132
 - training datasets, 129–130
- Deductive reasoning, 60
- Deep learning
 - convolutional neural
 - networks, 155
 - definition, 152–153
 - feedforward neural network, 154
 - radial basis neural
 - network, 156–157
- DeepMind WaveNet, 236
- Diabetes Digital Media
 - (DDM), 13, 323
- Diabetic foot concerns
 - (case study)
 - adoption strategy, 315
 - data governance, 312
 - data quality, 312
 - ethics governance, 313
 - hyperglycemia, 307
 - image recognition, 309
 - neural networks, 313
 - peripheral diabetic
 - neuropathy, 308
 - risks, 310
 - stakeholder understanding, 314
- Diagnosing disease risk (case study)
 - AI prioritization, 332
 - global diabetes population, 330
 - IoT, integration, 334
 - patient-data ecosystem, 328
 - predictive analytics, 333
 - real-world health statistics, 332
- Digicel Haiti network, 51
- Digital peer support, 208
- Digital surveys, 254
- 3-D printing, 208, 286
- Dimensionality reduction
 - algorithms, 162–164
- Dimension reduction techniques, 165
- Disease diagnosis, 96
- Distributed denial of service
 - (DDoS) attacks, 223–224
- Doctors and healthcare
 - professionals, AI
 - bioprinting and tissue
 - engineering, 287–288

- challenging decisions, 285
- diagnosing disease, 284
- 3-D printing, 286
- drug discovery, 285–286
- education, 288
- gene therapy, 289
- optimization, 284
- personalized prosthetics, 287
- pharmacology and
 - device, 288

Duplex software, 236

E

Economy and employment, 243–244

Electronic health record
(EHR), 45, 215

Emergency response team, 214

Ensemble

- bagging, 135
- boosting, 137–138
- random forests decision
 - tree, 136–137

Environment, machine

- learning, 80–81

Epilepsy (case study)

- antiepileptic drugs (AEDs), 321
- ketogenic diet (KD), 321
- KPE, 322
- randomized controlled trials
(RCTs), 322

Ethics

- AI models, 224
- machine learning, 224–225

Evaluation metrics

- classification
 - accuracy, 193
 - AUC, 195
 - confusion, 193–194
 - logarithmic loss, 195
 - per-class accuracy, 194
 - precision, recall, 196–197
- regression, 197–198

Evidence-based approach, 22

F

Fairness, 241

Feature, machine learning, 82

Foot ulcer detection, 208

Fraud detection, 96

Freedom of choice, 212–213

G

General Data Protection

- Regulation (GDPR), 217, 219, 252

Genetic algorithms (GA), 179, 181

Global standards, 247

Goals, machine learning, 79

H

Hadoop Distributed File System
(HDFS), 73

Healthcare

- evidence-based medicine, 261–264

INDEX

Healthcare (*cont.*)

- innovation *vs.* deliberation, 303–304
- personalized medicine, 264–266
- smart places, 300
 - smart homes, 301–302
 - smart hospitals, 302
- volume-to patient-centric value, 256–259, 261

Healthcare applications

- data-driven, patient, 12
- data privacy, 13
- diagnosis, 13
- drug discovery, 14
- machine learning, 12
- MRI scan, 12
- personalized treatment, 13
- prediction, 13

Healthcare data use cases

- electronic health records, 45–46
- evidence-based medicine, 49

IoT

- AFIB, 47
- innovative program, 48
- predictive analytics, 45
- public health, 50
- reduce readmission, 44
- waiting time, prediction, 44

Healthcare services, 237

Health intelligence, 237–238

Hippocratic Oath, data scientists, 253

Hold-back method, 106

Human bias, 226

Humanity, 234–235, 241–242, 248–249

Hyperparameter tuning

algorithm, 190, 200

grid search, 201

random search, 201

Hypoglycemia, 211–212

Hypothesis, machine

learning, 82

I, J

Identifiable data, 216

Individualized data, 216

Inductive reasoning, 60

Information governance, 246–247

Information security management

system (ISMS), 252

Informed consent, 212

Integration quality, 39

Intelligence, 235–236

Intelligence bias, 226–227

Interface agent, 78

International Organization for

Standardization (ISO), 67

Internet addiction, 242

Internet of Things (IoT), 5, 223

Iterative Dichotomizer 3

(ID3), 133

K

Ketogenic Program for Epilepsy

(KPE), 322

Key performance indicators

(KPIs), 44

Knowledge, machine learning, 80

L

Latent semantic analysis, 177

Learning method

applying, 103

test/training set, 105–106

choosing, 102

Liability

AI agent's behavior, 239

diagnosis of patient's

disease, 238–239

healthcare professionals, 238

organizational accountability,
239–240

Life and health insurance, 219, 237

Limited memory systems, 6

Linear regression, 139–141

Logistic regression, 141–142

M

Machine bias, 225–226

Machine learning, 8

algorithms, 43, 95–96, 218

method and results

assessment, 107

optimization, 113–114

ensembles, 115

implementation problem, 116

problem distribution, 115

supervised/unsupervised
learning, 115

principles, 83

problem specification, 97

categorization, 98

data error, 98

reporting results, 116–117

vs. software engineering, 84–85

Machine learning model

accuracy, 181, 184

experience (E), 120–123

libraries, 123–125

linearity, 185

parameters, 185–186

performance (P), 121, 181

productivity, 183

redundancy, 183

task (T), 120

Mapping function, 82

Master Data Management

(MDM), 69

Merged reality, 291

Metadata, 43

Monte-Carlo cross-validation, 107

Multivariate test

causation, 203

data drifts, 206

evaluation, 202

p value, 204–205

running, 205

statistical power, effect size, 204

use *t*-test, 204

MyFitnessPal, 220

N

Naïve Bayesian models,

145–148, 178

National Health Service

(NHS), 16

Natural language, 236

INDEX

Natural language processing
(NLP), [5](#), [24](#), [93](#), [167](#),
[169–170](#)

cosine similarity, [177–178](#)

lexical analysis, [170–171](#)

ngrams, [175–176](#)

Porter stemmer, [171](#)

TF-IDF vector, [176](#)

Tokenization, [169](#)

Netflix, [220](#)

Neural Networks

ANNs, [149](#)

perceptrons, [149–151](#)

n-fold cross-validation, [106](#)

Noninvasive blood glucose
levels, [207](#)

Nursing and delivery of
medicine, [292](#)

O

Open-source data
sharing, [208](#)

Organizational accountability

accuracy, [239](#)

auditability, [239](#)

explainability, [240](#)

fairness, [240](#)

responsibility, [240](#)

Organizations

code of conduct, [249](#)

data collection

laws and regulations, [252–253](#)

minimal amount, [251](#)

sensitive data, [252](#)

ethical code, [249](#), [251](#)

Overhype, [245](#)

P, Q

Pain management, [291](#)

Parametric *vs.* nonparametric
algorithms, [94](#)

Part of speech (POS) tag, [173–174](#)

Peer support communities, [208](#)

Performance, machine learning, [79](#)

Personalized medicine, [264–266](#)

Pharmacology, [222](#), [237](#)

Physical therapy, [292](#)

Playing God, [244](#)

Porter stemmer algorithm, [171–172](#)

Prediction ethics

agent learning, [232](#)

agile digital technologies, [230](#)

AI algorithms, [229](#)

healthcare, [233](#)

incorrect classification, [230](#)

information governance, [230](#)

intelligence, [230](#)

machine learning black
boxes, [229](#)

robots, [232](#)

technologies, [228](#)

Uber's Greyball algorithm, [231](#)

unintended consequences,
[233–234](#)

- validity, 231
- Volkswagen's algorithm, 231
- Predictive analytics, 223
 - patient monitoring, 57
 - personalized care, 57
 - techniques, 56
 - tools, 45
- Principal component analysis (PCA), 166
- Prioritizing treatments, 221
- Privacy, 220–221
- Psychometric and parametric testing, 218
- Public understanding, 214

R

- Random forest decision tree, 166
- Reactive systems, 6
- Real-time alerting, 49
- Real-world evidence, 222
- Recurrent Neural Networks (RNNs), 154–155
- Reductionism, 303
- Regulation, 245
- Reinforcement learning
 - algorithm, 90
 - reward-driven programming, 92
 - robotics, usage, 90
 - working, 91
- Robots, 298
 - companions, 299
 - drones, 299–300

- exoskeletons, 298
- inpatient care, 299
- robot-assisted surgery, 298
- Robustness, 45
- Root-mean-squared error (RMSE), 197

S

- Scaremongering, 245
- Screen time, 244
- Security, 223–224
- Semantic analysis, 175
- Semi-supervised learning, 90
- Sensitive data, 252
- Small data, 42
- Smartphones, 210
- Software as a service (SaaS), 33
- Speech recognition, 96
- Stakeholders, 245–246
- Staleness, 39
- Statistical model, 83
- Structured Query Language (SQL), 52
- Supervised and unsupervised models, 240
- Supervised learning
 - algorithm, 86, 125
 - classification, 127–128
 - filter methods, 126
 - regression, 128–129
 - classification, 87
 - forecasting, 87

INDEX

Supervised learning (*cont.*)

- logistic regression, [87](#)

- regression, [87](#)

- working, [86](#)

Support vector machine

- (SVM), [143–144](#)

Syntactic analysis, [172–173](#)

T

Technical governance, [245](#)

Technology addiction, [242–243](#)

Technology organizations, [254](#)

Telemedicine, [278](#)

Term Frequency-Inverse

- Document Frequency

- (TF-IDF) vector, [176](#)

Theory of mind, [7](#)

Tidy data, [101](#)

Time series analysis techniques, [44](#)

Training data, [81](#)

Treatments and management

- pathways, [222](#)

Try many algorithms, [107](#)

Type 2 diabetes (case

- study), [240, 248](#)

- carbohydrate-reduced diet, [316](#)

- low-carb program (LCP), [317](#)

- quasi-experimental

- research, [317](#)

U

Uber's Greyball algorithm, [231](#)

Unintended consequences,

- [233–234](#)

Unsupervised learning, [157](#)

- association, [89](#)

- clustering, [89, 158](#)

- dimensionality reduction, [89](#)

- feature extraction methods, [89](#)

- KMeans, [158, 160](#)

- working, [88](#)

Utility, machine learning, [79](#)

V, W, X, Y, Z

Validation, machine learning, [82](#)

Variance, training set, [110](#)

Veracity

- clean data, [39](#)

- collaborative data, [40](#)

- complaint data, [40](#)

- consistent data, [40](#)

- consolidated data, [40](#)

- current data, [40](#)

Virtual appointments and

- classrooms, [293](#)

Virtual assistants, [241](#)

Virtual reality (VR), [290](#)

Volkswagen's algorithm, [231](#)