

Building Intelligent Systems

**A Guide to Machine Learning
Engineering**

Geoff Hulten

Apress®

Building Intelligent Systems: A Guide to Machine Learning Engineering

Geoff Hulten

Lynnwood, Washington, USA

ISBN-13 (pbk): 978-1-4842-3431-0

ISBN-13 (electronic): 978-1-4842-3432-7

<https://doi.org/10.1007/978-1-4842-3432-7>

Library of Congress Control Number: 2018934680

Copyright © 2018 by Geoff Hulten

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: Susan McDermott

Development Editor: Laura Berendson

Coordinating Editor: Rita Fernando

Cover designed by eStudioCalamar

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

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/9781484234310. For more detailed information, please visit <http://www.apress.com/source-code>.

Printed on acid-free paper

To Dad, for telling me what I needed to hear.
To Mom, for pretty much just telling me what I wanted to hear.
And to Nicole.

Table of Contents

About the Author	xvii
About the Technical Reviewer	xix
Acknowledgments	xxi
Introduction	xxiii
Part I: Approaching an Intelligent Systems Project.....	1
Chapter 1: Introducing Intelligent Systems	3
Elements of an Intelligent System	4
An Example Intelligent System	6
The Internet Toaster.....	6
Using Data to Toast.....	7
Sensors and Heuristic Intelligence.....	8
Toasting with Machine Learning.....	10
Making an Intelligent System	11
Summary.....	12
For Thought.....	13
Chapter 2: Knowing When to Use Intelligent Systems.....	15
Types of Problems That Need Intelligent Systems	15
Big Problems	16
Open-Ended Problems.....	16
Time-Changing Problems	17
Intrinsically Hard Problems	18
Situations When Intelligent Systems Work.....	18
When a Partial System Is Viable and Interesting	19
When You Can Use Data from the System to Improve	19

TABLE OF CONTENTS

- When the System Can Interface with the Objective 20
- When it is Cost Effective..... 21
- When You Aren't Sure You Need an Intelligent System..... 22
- Summary..... 23
- For Thought..... 23
- Chapter 3: A Brief Refresher on Working with Data 25**
- Structured Data..... 25
- Asking Simple Questions of Data..... 27
- Working with Data Models 29
- Conceptual Machine Learning 30
- Common Pitfalls of Working with Data 31
- Summary..... 33
- For Thought..... 34
- Chapter 4: Defining the Intelligent System's Goals 35**
- Criteria for a Good Goal..... 36
- An Example of Why Choosing Goals Is Hard 36
- Types of Goals 38
- Organizational Objectives..... 38
- Leading Indicators..... 39
- User Outcomes 41
- Model Properties 42
- Layering Goals 43
- Ways to Measure Goals..... 44
- Waiting for More Information..... 44
- A/B Testing..... 45
- Hand Labeling..... 45
- Asking Users..... 46
- Decoupling Goals..... 46

Keeping Goals Healthy	47
Summary.....	48
For Thought.....	48
Part II: Intelligent Experiences	51
Chapter 5: The Components of Intelligent Experiences	53
Presenting Intelligence to Users	54
An Example of Presenting Intelligence	55
Achieve the System's Objectives	57
An Example of Achieving Objectives	58
Minimize Intelligence Flaws.....	58
Create Data to Grow the System.....	59
An Example of Collecting Data	60
Summary.....	61
For Thought.....	62
Chapter 6: Why Creating Intelligent Experiences Is Hard	63
Intelligence Make Mistakes	63
Intelligence Makes Crazy Mistakes.....	65
Intelligence Makes Different Types of Mistakes.....	66
Intelligence Changes.....	68
The Human Factor.....	70
Summary.....	72
For Thought.....	72
Chapter 7: Balancing Intelligent Experiences.....	75
Forcefulness	76
Frequency	78
Value of Success.....	79
Cost of Mistakes	81
Knowing There Is a Mistake	81
Recovering from a Mistake.....	82

TABLE OF CONTENTS

- Intelligence Quality 83
- Summary..... 85
- For Thought..... 86
- Chapter 8: Modes of Intelligent Interaction 87**
- Automate..... 87
- Prompt 89
- Organize..... 90
- Annotate..... 92
- Hybrid Experiences 93
- Summary..... 94
- For Thought..... 95
- Chapter 9: Getting Data from Experience 97**
- An Example: TeamMaker..... 98
 - Simple Interactions..... 98
 - Making It Fun..... 99
 - Connecting to Outcomes 100
- Properties of Good Data 100
 - Context, Actions, and Outcomes 101
 - Good Coverage 102
 - Real Usage..... 103
 - Unbiased..... 103
 - Does Not Contain Feedback Loops..... 104
 - Scale..... 105
- Ways to Understand Outcomes 106
 - Implicit Outcomes..... 106
 - Ratings 107
 - Reports 107
 - Escalations 108
 - User Classifications..... 108
- Summary..... 109
- For Thought..... 110

Chapter 10: Verifying Intelligent Experiences	111
Getting Intended Experiences	112
Working with Context	112
Working with Intelligence	114
Bringing it Together	115
Achieving Goals.....	116
Continual Verification	117
Summary.....	117
For Thought	118
Part III: Implementing Intelligence	121
Chapter 11: The Components of an Intelligence Implementation.....	123
An Example of Intelligence Implementation.....	124
Components of an Intelligence Implementation	127
The Intelligence Runtime.....	127
Intelligence Management.....	127
Intelligence Telemetry Pipeline.....	128
The Intelligence Creation Environment.....	128
Intelligence Orchestration	129
Summary.....	130
For Thought	130
Chapter 12: The Intelligence Runtime	133
Context.....	134
Feature Extraction.....	135
Models	137
Execution	138
Results	139
Instability in Intelligence.....	139
Intelligence APIs	140
Summary.....	140
For Thought	141

TABLE OF CONTENTS

- Chapter 13: Where Intelligence Lives 143**
 - Considerations for Positioning Intelligence..... 143
 - Latency in Updating..... 144
 - Latency in Execution..... 146
 - Cost of Operation..... 148
 - Offline Operation..... 149
 - Places to Put Intelligence..... 150
 - Static Intelligence in the Product 150
 - Client-Side Intelligence 151
 - Server-Centric Intelligence..... 152
 - Back-End (Cached) Intelligence..... 153
 - Hybrid Intelligence..... 154
 - Summary..... 155
 - For Thought..... 156

- Chapter 14: Intelligence Management..... 157**
 - Overview of Intelligence Management 157
 - Complexity in Intelligent Management 158
 - Frequency in Intelligence Management 159
 - Human Systems 159
 - Sanity-Checking Intelligence 160
 - Checking for Compatibility 160
 - Checking for Runtime Constraints..... 161
 - Checking for Obvious Mistakes 162
 - Lighting Up Intelligence 162
 - Single Deployment 163
 - Silent Intelligence..... 164
 - Controlled Rollout..... 165
 - Flighting..... 166
 - Turning Off Intelligence..... 167
 - Summary..... 168
 - For Thought..... 168

Chapter 15: Intelligent Telemetry	171
Why Telemetry Is Needed.....	171
Make Sure Things Are Working.....	172
Understand Outcomes	173
Gather Data to Grow Intelligence.....	174
Properties of an Effective Telemetry System	175
Sampling	175
Summarizing	176
Flexible Targeting	177
Common Challenges	178
Bias.....	178
Rare Events	179
Indirect Value.....	180
Privacy.....	180
Summary.....	181
For Thought	182
Part IV: Creating Intelligence.....	183
Chapter 16: Overview of Intelligence	185
An Example Intelligence.....	185
Contexts	187
Implemented at Runtime	187
Available for Intelligence Creation.....	189
Things Intelligence Can Predict	190
Classifications	190
Probability Estimates.....	191
Regressions.....	193
Rankings.....	194
Hybrids and Combinations.....	194
Summary.....	194
For Thought	195

TABLE OF CONTENTS

- Chapter 17: Representing Intelligence 197**
 - Criteria for Representing Intelligence 197
 - Representing Intelligence with Code 198
 - Representing Intelligence with Lookup Tables..... 199
 - Representing Intelligence with Models..... 201
 - Linear Models..... 202
 - Decision Trees 203
 - Neural Networks..... 205
 - Summary..... 207
 - For Thought..... 207
- Chapter 18: The Intelligence Creation Process..... 209**
 - An Example of Intelligence Creation: Blinker 210
 - Understanding the Environment 210
 - Define Success 212
 - Get Data 213
 - Bootstrap Data..... 214
 - Data from Usage..... 215
 - Get Ready to Evaluate 216
 - Simple Heuristics..... 217
 - Machine Learning 218
 - Understanding the Tradeoffs..... 219
 - Assess and Iterate 219
 - Maturity in Intelligence Creation..... 220
 - Being Excellent at Intelligence Creation 221
 - Data Debugging 221
 - Verification-Based Approach 222
 - Intuition with the Toolbox 222
 - Math (?) 223
 - Summary..... 223
 - For Thought..... 224

Chapter 19: Evaluating Intelligence.....	225
Evaluating Accuracy	226
Generalization.....	226
Types of Mistakes.....	227
Distribution of Mistakes.....	230
Evaluating Other Types of Predictions.....	230
Evaluating Regressions	230
Evaluating Probabilities.....	231
Evaluating Rankings	231
Using Data for Evaluation.....	232
Independent Evaluation Data.....	232
Independence in Practice	233
Evaluating for Sub-Populations	235
The Right Amount of Data.....	237
Comparing Intelligences	238
Operating Points	238
Curves	239
Subjective Evaluations	240
Exploring the Mistakes	241
Imagining the User Experience.....	242
Finding the Worst Thing.....	242
Summary.....	243
For Thought	244
Chapter 20: Machine Learning Intelligence.....	245
How Machine Learning Works	245
The Pros and Cons of Complexity	247
Underfitting.....	248
Overfitting.....	249
Balancing Complexity	249
Feature Engineering.....	250
Converting Data to Useable Format.....	251

TABLE OF CONTENTS

- Helping your Model Use the Data 253
- Normalizing 254
- Exposing Hidden Information 255
- Expanding the Context..... 256
- Eliminating Misleading Things..... 256
- Modeling 257
 - Complexity Parameters 258
 - Identifying Overfitting 259
- Summary..... 260
- For Thought..... 261
- Chapter 21: Organizing Intelligence 263**
 - Reasons to Organize Intelligence..... 263
 - Properties of a Well-Organized Intelligence..... 264
 - Ways to Organize Intelligence..... 265
 - Decouple Feature Engineering 266
 - Multiple Model Searches..... 268
 - Chase Mistakes 269
 - Meta-Models 270
 - Model Sequencing..... 272
 - Partition Contexts 274
 - Overrides 275
 - Summary..... 277
 - For Thought..... 278
- Part V: Orchestrating Intelligent Systems 279**
- Chapter 22: Overview of Intelligence Orchestration..... 281**
 - Properties of a Well-Orchestrated Intelligence 282
 - Why Orchestration Is Needed..... 282
 - Objective Changes..... 283
 - Users Change 284
 - Problem Changes 285

Intelligence Changes	286
Costs Change.....	287
Abuse.....	287
The Orchestration Team	288
Summary.....	288
For Thought.....	289
Chapter 23: The Intelligence Orchestration Environment.....	291
Monitor the Success Criteria.....	292
Inspect Interactions	293
Balance the Experience	295
Override Intelligence	296
Create Intelligence	298
Summary.....	299
For Thought.....	300
Chapter 24: Dealing with Mistakes.....	301
The Worst Thing That Could Happen	301
Ways Intelligence Can Break.....	303
System Outage	303
Model Outage	304
Intelligence Errors	304
Intelligence Degradation	305
Mitigating Mistakes	306
Invest in Intelligence	306
Balance the Experience.....	307
Adjust Intelligence Management Parameters.....	307
Implement Guardrails	308
Override Errors	308
Summary.....	309
For Thought.....	310

TABLE OF CONTENTS

- Chapter 25: Adversaries and Abuse 311**
 - Abuse Is a Business 312
 - Abuse Scales 313
 - Estimating Your Risk..... 313
 - What an Abuse Problem Looks Like..... 314
 - Ways to Combat Abuse 315
 - Add Costs 315
 - Becoming Less Interesting to Abusers 315
 - Machine Learning with an Adversary 316
 - Get the Abuser out of the Loop..... 316
 - Summary..... 316
 - For Thought..... 317
- Chapter 26: Approaching Your Own Intelligent System..... 319**
 - An Intelligent System Checklist 319
 - Approach the Intelligent System Project 320
 - Plan for the Intelligent Experience 321
 - Plan the Intelligent System Implementation..... 323
 - Get Ready to Create Intelligence 325
 - Orchestrate Your Intelligent System 327
 - Summary..... 329
 - For Thought..... 329
- Index..... 331**

About the Author



Geoff Hulten is a machine learning scientist and PhD in machine learning. He has managed applied machine learning teams for over a decade, building dozens of Internet-scale Intelligent Systems that have hundreds of millions of interactions with users every day. His research has appeared in top international conferences, received thousands of citations, and won a SIGKDD Test of Time award for influential contributions to the data mining research community that have stood the test of time.

About the Technical Reviewer



Jeb Haber has a BS in Computer Science from Willamette University. He spent nearly two decades at Microsoft working on a variety of projects across Windows, Internet Explorer, Office, and MSN. For the last decade-plus of his Microsoft career, Jeb led the program management team responsible for the safety and security services provided by Microsoft SmartScreen (anti-phishing, anti-malware, and so on.) Jeb's team developed and managed global-scale Intelligent Systems with hundreds of millions of users. His role included product vision/planning/strategy, project management, metrics definition and people/team development. Jeb

helped organize a culture along with the systems and processes required to repeatedly build and run global scale, 24×7 intelligence and reputation systems. Jeb is currently serving as the president of two non-profit boards for organizations dedicated to individuals and families dealing with the rare genetic disorder phenylketonuria (PKU).

Acknowledgments

There are so many people who were part of the Intelligent Systems I worked on over the years. These people helped me learn, helped me understand. In particular, I'd like to thank:

Jeb Haber and John Scarrow for being two of the key minds in developing the concepts described in this book and for being great collaborators over the years. None of this would have happened without their leadership and dedication.

Also: Anthony P., Tomasz K., Rob S., Rob M., Dave D., Kyle K., Eric R., Ameya B., Kris I., Jeff M., Mike C., Shankar S., Robert R., Chris J., Susan H., Ivan O., Chad M. and many others...

Introduction

Building Intelligent Systems is a book about leveraging machine learning in practice.

It covers everything you need to produce a fully functioning Intelligent System, one that leverages machine learning and data from user interactions to improve over time and achieve success.

After reading this book you'll be able to design an Intelligent System end-to-end.

You'll know:

- When to use an Intelligent System and how to make it achieve your goals.
- How to design effective interactions between users and Intelligent Systems.
- How to implement an Intelligent System across client, service, and back end.
- How to build the intelligence that powers an Intelligent System and grow it over time.
- How to orchestrate an Intelligent System over its life-cycle.

You'll also understand how to apply your existing skills, whether in software engineering, data science, machine learning, management or program management to the effort.

There are many great books that teach data and machine-learning skills. Those books are similar to books on programming languages; they teach valuable skills in great detail. This book is more like a book on software engineering; it teaches how to take those base skills and produce working systems.

This book is based on more than a decade of experience building Internet-scale Intelligent Systems that have hundreds of millions of user interactions per day in some of the largest and most important software systems in the world. I hope this book helps accelerate the proliferation of systems that turn data into impact and helps readers develop practical skills in this important area.

Who This Book Is For

This book is for anyone with a computer science degree who wants to understand what it takes to build effective Intelligent Systems.

Imagine a typical software engineer who is assigned to a machine learning project. They want to learn more about it so they pick up a book, and it is technical, full of statistics and math and modeling methods. These are important skills, but they are the wrong information to help the software engineer contribute to the effort. *Building Intelligent Systems* is the right book for them.

Imagine a machine learning practitioner who needs to understand how the end-to-end system will interact with the models they produce, what they can count on, and what they need to look out for in practice. *Building Intelligent Systems* is the right book for them.

Imagine a technical manager who wants to begin benefiting from machine learning. Maybe they hire a machine learning PhD and let them work for a while. The machine learning practitioner comes back with charts, precision/recall curves, and training data requests, but no framework for how they should be applied. *Building Intelligent Systems* is the right book for that manager.

Data and Machine Learning Practitioners

Data and machine learning are at the core of many Intelligent Systems, but there is an incredible amount of work to be done between the development of a working model (created with machine learning) and the eventual sustainable customer impact. Understanding this supporting work will help you be better at modeling in a number of ways.

First, it's important to **understand the constraints** these systems put on your modeling. For example, where will the model run? What data will it have access to? How fast does it need to be? What is the business impact of a false positive? A false negative? How should the model be tuned to maximize business results?

Second, it's important to be able to **influence the other participants**. Understanding the pressures on the engineers and business owners will help you come to good solutions and maximize your chance for success. For example, you may not be getting all the training data you'd like because of telemetry sampling. Should you double down on

modeling around the problem, or would an engineering solution make more sense? Or maybe you are being pushed to optimize for a difficult extremely-high precision, when your models are already performing at a very good (but slightly lower) precision. Should you keep chasing that super-high precision or should you work to influence the user experience in ways that reduce the customer impact of mistakes?

Third, it's important to understand how the **supporting systems can benefit you**. The escalation paths, the manual over-rides, the telemetry, the guardrails that prevent against major mistakes—these are all tools you can leverage. You need to understand when to use them and how to integrate them with your modeling process. Should you discard a model that works acceptably for 99% of users but really, really badly for 1% of users? Or maybe you can count on other parts of the system to address the problem.

Software Engineers

Building software that delights customers is a lot of work. No way around it, behind every successful software product and service there is some serious engineering. Intelligent Systems have some unique properties which present interesting challenges. This book describes the associated concepts so you can design and build Intelligent Systems that are efficient, reliable, and that best-unlock the power of machine learning and data science.

First, this book will identify the **entities and abstractions** that need to exist within a successful Intelligent System. You will learn the concepts behind the intelligence runtime, context and features, models, telemetry, training data, intelligence management, orchestration, and more.

Second, the book will give you a **conceptual understanding of machine learning and data sciences**. These will prepare you to have good discussions about tradeoffs between engineering investments and modeling investments. Where can a little bit of your work really enable a solution? And where are you being asked to boil the ocean to save a little bit of modeling time?

Third, the book will explore **patterns for Intelligent Systems** that my colleagues and I have developed over a decade and through implementing many working systems. What are the pros and cons of running intelligence in a client or in a service? How do you bound and verify components that are probabilistic? What do you need to include in telemetry so the system can evolve?

Program Managers

Machine learning and Data Sciences are hot topics. They are fantastic tools, but they are tools; they are not solutions. This book will give you enough conceptual understanding so you know what these tools are good at and how to deploy them to solve your business problems.

The first thing you'll learn is to develop an **intuition for when machine learning and data science are appropriate**. There is nothing worse than trying to hammer a square peg into a round hole. You need to understand what types of problems can be solved by machine learning. But just as importantly, you need to understand what types of problems can't be—or at least not easily. There are so many participants in a successful endeavor, and they speak such different, highly-technical, languages, that this is particularly difficult. This book will help you understand enough so you can ask the right questions and understand what you need from the answers.

The second is to get an intuition on return on investment so you can determine **how much Intelligent System to use**. By understanding the real costs of building and maintaining a system that turns data into impact you can make better choices about when to do it. You can also go into it with open eyes, and have the investment level scoped for success. Sometimes you need all the elements described in this book, but sometimes the right choice for your business is something simpler. This book will help you make good decisions and communicate them with confidence and credibility.

Finally, the third thing a program manager will learn here is to understand how to **plan, staff, and manage an Intelligent System project**. You will get the benefit of our experience building many large-scale Intelligent Systems: the life cycle of an Intelligent System; the day-to-day process of running it; the team and skills you need to succeed.