



Artificial Intelligence for Marketing





Artificial Intelligence for Marketing

Practical Applications



Jim Sterne



WILEY





Copyright © 2017 by Rising Media, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.
Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at www.wiley.com/go/permissions.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993, or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Cataloging-in-Publication Data is Available:

ISBN 9781119406334 (Hardcover)

ISBN 9781119406372 (ePDF)

ISBN 9781119406365 (ePub)

Cover Design: Wiley

Cover Image: © Kngkyle2/Getty Images

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1





Contents

[Foreword by Tom Davenport](#) [xiii](#)

[Preface](#) [xvii](#)

[Acknowledgments](#) [xix](#)

[Chapter 1 Welcome to the Future](#) [1](#)

[Welcome to Autonomic Marketing](#) [3](#)

[Welcome to Artificial Intelligence for Marketers](#) [3](#)

[Whom Is This Book For?](#) [5](#)

[The Bright, Bright Future](#) [6](#)

[Is AI So Great if It's So Expensive?](#) [7](#)

[What's All This AI Then?](#) [9](#)

[The AI Umbrella](#) [9](#)

[The Machine that Learns](#) [10](#)

[Are We There Yet?](#) [14](#)

[AI-pocalypse](#) [15](#)

[Machine Learning's Biggest Roadblock](#) [23](#)

[Machine Learning's Greatest Asset](#) [24](#)

[Are We Really Calculable?](#) [56](#)

[Chapter 2 Introduction to Machine Learning](#) [59](#)

[Three Reasons Data Scientists Should Read This Chapter](#) [59](#)

[Every Reason Marketing Professionals Should Read](#)

[This Chapter](#) [60](#)

[We Think We're So Smart](#) [60](#)

[Define Your Terms](#) [61](#)

[All Models Are Wrong](#) [62](#)

[Useful Models](#) [64](#)

[Too Much to Think About](#) [66](#)

[Machines Are Big Babies](#) [68](#)

[Where Machines Shine](#) [69](#)

[Strong versus Weak AI](#) [71](#)

[The Right Tool for the Right Job](#) [72](#)

[Make Up Your Mind](#) [88](#)

[One Algorithm to Rule Them All?](#) [89](#)

[Accepting Randomness](#) [92](#)



[X ▶ CONTENTS](#)

[Which Tech Is Best? 94](#)
[For the More Statistically Minded 94](#)
[What Did We Learn? 101](#)

[Chapter 3 Solving the Marketing Problem 103](#)

[One-to-One Marketing 105](#)
[One-to-Many Advertising 107](#)
[The Four Ps 108](#)
[What Keeps a Marketing Professional Awake? 109](#)
[The Customer Journey 111](#)
[We Will Never Really Know 111](#)
[How Do I Connect? Let Me Count the Ways 114](#)
[Why Do I Connect? Branding 117](#)
[Marketing Mix Modeling 119](#)
[Econometrics 121](#)
[Customer Lifetime Value 121](#)
[One-to-One Marketing—The Meme 122](#)
[Seat-of-the-Pants Marketing 123](#)
[Marketing in a Nutshell 124](#)
[What Seems to Be the Problem? 126](#)

[Chapter 4 Using AI to Get Their Attention 128](#)

[Market Research: Whom Are We After? 128](#)
[Marketplace Segmentation 131](#)
[Raising Awareness 141](#)
[Social Media Engagement 155](#)
[In Real Life 158](#)
[The B2B World 158](#)

[Chapter 5 Using AI to Persuade 165](#)

[The In-Store Experience 168](#)
[On the Phone 178](#)
[The Onsite Experience—Web Analytics 179](#)
[Merchandising 186](#)
[Closing the Deal 188](#)
[Back to the Beginning: Attribution 193](#)

[Chapter 6 Using AI for Retention 200](#)

[Growing Customer Expectations 200](#)
[Retention and Churn 202](#)
[Many Unhappy Returns 204](#)
[Customer Sentiment 208](#)
[Customer Service 209](#)
[Predictive Customer Service 216](#)



[Chapter 7 The AI Marketing Platform 218](#)[Supplemental AI 218](#)[Marketing Tools from Scratch 221](#)[A Word about Watson 224](#)[Building Your Own 230](#)[Chapter 8 Where Machines Fail 232](#)[A Hammer Is Not a Carpenter 232](#)[Machine Mistakes 235](#)[Human Mistakes 241](#)[The Ethics of AI 247](#)[Solution? 258](#)[What Machines Haven't Learned Yet 260](#)[Chapter 9 Your Strategic Role in Onboarding AI 262](#)[Getting Started, Looking Forward 264](#)[AI to Leverage Humans 272](#)[Collaboration at Work 274](#)[Your Role as Manager 276](#)[Know Your Place 282](#)[AI for Best Practices 286](#)[Chapter 10 Mentoring the Machine 289](#)[How to Train a Dragon 290](#)[What Problem Are You Trying to Solve? 291](#)[What Makes a Good Hypothesis? 294](#)[The Human Advantage 297](#)[Chapter 11 What Tomorrow May Bring 305](#)[The Path to the Future 307](#)[Machine, Train Thyself 308](#)[Intellectual Capacity as a Service 308](#)[Data as a Competitive Advantage 310](#)[How Far Will Machines Go? 316](#)[Your Bot Is Your Brand 319](#)[My AI Will Call Your AI 321](#)[Computing Tomorrow 325](#)[About the Author 327](#)[Index 329](#)





Foreword

Thomas H. Davenport

*Distinguished Professor, Babson College and Research Fellow, MIT
Author of *Competing on Analytics* and *Only Humans Need Apply**

Forewords to books can play a variety of roles. One is to describe in more general terms what the book is about. That's not really necessary, since Jim Sterne is a master at communicating complex topics in relatively simple terms.

Another common purpose is to describe how the book fits into the broader literature on the topic. That doesn't seem necessary in this case, either, since there isn't much literature on artificial intelligence (AI) for marketing, and even if there were, you've probably turned to this book to get one easy-to-consume source.

A third possible objective for forewords is to persuade you of the importance and relevance of the book, with the short-term goal of having you actually buy it or read onward if you already bought it. I'll adopt that goal, and provide external testimony that AI already is important to marketing, that it will become much more so in the future, and that any good marketing executive needs to know what it can do.

It's not that difficult to argue that marketing in the future will make increasing use of AI. Even today, the components of an AI-based approach are largely in place. Contemporary marketing is increasingly quantitative, targeted, and tied to business outcomes. Ads and promotions are increasingly customized to individual consumers in real time. Companies employ multiple channels to get to customers, but all of them increasingly employ digital content. Company marketers still work with agencies, many of which have developed analytical capabilities of their own.

As Sterne points out, data is the primary asset for AI-based marketing approaches. Data for marketing comes from a company's own systems, agencies, third-party syndicators, customer online behaviors, and many other sources—and certainly comprises “big data” in the aggregate. About 25 percent of today's marketing budgets are devoted to digital channels, and almost 80 percent of marketing organizations make technology-oriented capital expenditures—typically hardware and software—according to a recent Gartner survey. Clearly some of that capital will be spent on AI.





xiv ► FOREWORD

Companies still try to maintain a consistent brand image, but the annual marketing strategy is increasingly a relic of the past. Instead of making a few major decisions each year, companies or their agencies make literally thousands of real-time decisions a day about which ads to run on which sites, which search terms to buy, which version of a website to adopt, and so forth. Even the choice of what service providers and marketing software vendors to work with is complex enough to deserve a decision-making algorithm.

Already there are simply too many decisions involving too many complex variables and too much data for humans to make all of them. Marketing activities and decisions are increasing far more rapidly than marketing budgets or the numbers and capabilities of human marketers. An increasing number of marketing decisions employ some sort of AI, and this trend will only increase.

Companies are typically trying to define and target specific customers or segments, and if there are thousands or millions of customers, AI is needed to get to that level of detail. Companies also want to customize the experience of the customer, and that also requires machine learning or some other form of AI. AI can also help to deliver value across omnichannel customer relationships, and to ensure effective communications at all customer touchpoints. Finally, AI can help companies make decisions with similar criteria across the digital and analog marketing worlds.

Today, AI in marketing supports only certain kinds of decisions. They are typically repetitive decisions based on data, and each decision has low monetary value (though in total they add up to large numbers). AI-based decisions today primarily involve digital content and channels or online promotions. Of course, almost all content is becoming digitized, so it makes for a pretty big category. This set of AI-supported activities includes digital advertising buys (called *programmatic buying*), website operation and optimization, search engine optimization, A/B testing, outbound e-mail marketing, lead filtering and scoring, and many other marketing tasks.

And it seems highly likely that this list will continue to grow. Television advertising—the mainstay of large companies' marketing activities for many years—is moving toward a programmatic buying model. Creative brand development activities are still largely done by humans, but the decisions about which images and copy will be adopted are now sometimes made through AI-based testing. High-level decisions about marketing mix and resource allocation are still ultimately made by marketing executives, but they are usually done with software and are often performed more frequently than annually.





Preface

If you're in marketing, AI is a powerful ally.
If you're in data science, marketing is a rich problem set.

Artificial Intelligence (AI) had a breakthrough year in 2016, not only with machine learning, but with public awareness as well. And it's only going to continue. This year, most marketers believe consumers are ready for the technology.

"Artificial Intelligence Roundup," eMarketer, February 2017

AI IN A NUTSHELL

Artificial intelligence (AI) is the next, logical step in computing: a program that can figure out things for itself. It's a program that can reprogram itself.

The Three Ds of Artificial Intelligence

The shorthand for remembering what's special about AI is that it can *detect*, *deliberate*, and *develop*—all on its own.

Detect

Artificial intelligence can discover which elements or attributes in a bunch of data are the most predictive. Even when there is a massive amount of data made up of lots of different *kinds* of data, AI can identify the most revealing characteristics, figuring out which to pay attention to and which to ignore.

Deliberate

AI can infer rules about the data, *from* that data, and weigh the most predictive attributes against each other to answer a question or make a recommendation. It can ponder the relevance of each and reach a conclusion.





Develop

AI can grow and mature with each iteration. It can alter its opinion about the environment as well as how it evaluates that environment based on new information or the results of experimentation. It can program itself.

An individual's search terms are more important than her location, which is more important than her age (detect). When people use six or more words in a search, their propensity to purchase is so high that a discount is counterproductive (deliberate). Once it is noted that women under the age of 24 are not likely to purchase, regardless of words in a search, an experiment can be run to offer them free shipping (develop).

THIS IS YOUR MARKETING ON AI

The tools are not supernatural. They are not beyond the understanding of mortals. You owe it to yourself to understand how they are about to rock your world.

Intelligence is the ability to adapt to change.

—Stephen Hawking

The companion website for *Artificial Intelligence for Marketing: Practical Applications* can be found at: AI4Marketing.com.





Acknowledgments

I am forever grateful to the many people who have blogged, tweeted, published videos on, and answered my questions about artificial intelligence and machine learning.

Specifically, thanks go to Barry Levine, Bob Page, Brent Dykes, Brian Solis, Christopher Berry, Dan McCarthy, Dave Smith, David Raab, Dean Abbott, Dennis Mortensen, Doc Searls, Eric Siegel, Gary Angel, Himanshu Sharma, Ian Thomas, Kaj van de Loo, Mark Gibbs, Matt Gershoff, Matthew Todd, Michael Rappa, Michael Wu, Michelle Street, Pat LaPointe, Peter Fader, Rohit Rudrapatna, Ron Kohavi, Russ Klein, Russell McAthy, Scott Brinker, Scott Litman, Tim Wilson, Tom Cunniff, Tom Davenport, Tom Mitchell, Tyler Vigen, Vicky Brock, and Vincent Granville.

And, as always, Matt Cutler.







Artificial Intelligence for Marketing





2 ▶ ARTIFICIAL INTELLIGENCE FOR MARKETING

- The Ford Motor Company produced its first car—the Ford Model A.
- Thomas Edison invented the nickel-alkaline storage battery.
- The first electric typewriter was invented by George Canfield Blickensderfer of Erie, Pennsylvania.
- The first radio that successfully received a radio transmission was developed by Guglielmo Marconi.
- The Wright brothers flew at Kitty Hawk.
- The Panama Canal was under construction.
- Benjamin Holt invented one of the first practical continuous tracks for use in tractors and tanks.
- The Victor Talking Machine Company released the Victrola.
- The Autochrome Lumière, patented in 1903, became the first commercial color photography process.

My grandfather then lived to see men walk on the moon.

In the next few decades, we will see:

- Self-driving cars replace personally owned transportation.
- Doctors routinely operate remote, robotic surgery devices.
- Implantable communication devices replace mobile phones.
- In-eye augmented reality become normalized.
- Maglev elevators travel sideways and transform building shapes.
- Every surface consume light for energy and act as a display.
- Mind-controlled prosthetics with tactile skin interfaces become mainstream.
- Quantum computing make today's systems microscopic.
- 3-D printers allow for instant delivery of goods.
- Style-selective, nanotech clothing continuously clean itself.

And today's youngsters will live to see a colony on Mars.

It's no surprise that computational systems will manage more tasks in advertising and marketing. Yes, we have lots of technology for marketing, but the next step into artificial intelligence and machine learning will be different. Rather than being an ever-larger confusion of rules-based programs, operating faster than the eye can see, AI systems will operate more inscrutably than the human mind can fathom.





WELCOME TO AUTONOMIC MARKETING

The autonomic nervous system controls everything you don't have to think about: your heart, your breathing, your digestion. All of these things can happen while you're asleep or unconscious. These tasks are complex, interrelated, and vital. They are so necessary they must function continuously without the need for deliberate thought.

That's where marketing is headed. We are on the verge of the need for autonomic responses just to stay afloat. Personalization, recommendations, dynamic content selection, and dynamic display styles are all going to be table stakes.

The technologies seeing the light of day in the second decade of the twenty-first century will be made available as services and any company *not* using them will suffer the same fate as those that decided not to avail themselves of word processing, database management, or Internet marketing. And so, it's time to open up that black box full of mumbo-jumbo called artificial intelligence and understand it just well enough to make the most of it for marketing. Ignorance is no excuse. You should be comfortable enough with artificial intelligence to put it to practical use without having to get a degree in data science.

WELCOME TO ARTIFICIAL INTELLIGENCE FOR MARKETERS

It is of the highest importance in the art of detection to be able to recognize, out of a number of facts, which are incidental and which vital.

Sherlock Holmes, The Reigate Squires

This book looks at some current buzzwords to make just enough sense for regular marketing folk to understand what's going on.

- This is no deep exposé on the dark arts of artificial intelligence.
- This is no textbook for learning a new type of programming.
- This is no exhaustive catalog of cutting-edge technologies.

This book is not for those with advanced math degrees or those who wish to become data scientists. If, however, you are inspired to delve into the bottomless realm of modern systems building, I'll point you to "How to Get the Best Deep Learning Education for Free"² and be happy to take the credit for inspiring you. But that is not my intent.





4 ▶ ARTIFICIAL INTELLIGENCE FOR MARKETING

You *will not* find passages like the following in this book:

Monte-Carlo simulations are used in many contexts: to produce high quality pseudo-random numbers, in complex settings such as multi-layer spatio-temporal hierarchical Bayesian models, to estimate parameters, to compute statistics associated with very rare events, or even to generate large amount of data (for instance cross and auto-correlated time series) to test and compare various algorithms, especially for stock trading or in engineering.

“24 Uses of Statistical Modeling” (Part II)³

You *will* find explanations such as: Artificial intelligence is valuable because it was designed to deal in gray areas rather than crank out statistical charts and graphs. It is capable, over time, of understanding context.

The purpose of this tome is to be a primer, an introduction, a statement of understanding for those who have regular jobs in marketing—and would like to keep them in the foreseeable future.

Let’s start with a super-simple comparison between artificial intelligence and machine learning from Avinash Kaushik, digital marketing evangelist at Google: “AI is an *intelligent machine* and ML is the *ability to learn without being explicitly programmed*.”

Artificial intelligence is a machine pretending to be a human. Machine learning is a machine pretending to be a statistical programmer. Managing either one requires a data scientist.

An ever-so-slightly deeper definition comes from E. Fredkin University professor at the Carnegie Mellon University Tom Mitchell:⁴

The field of Machine Learning seeks to answer the question, “How can we build computer systems that automatically improve with experience, and what are the fundamental laws that govern all learning processes?”

A machine learns with respect to a particular task T, performance metric P, and type of experience E, if the system reliably improves its performance P at task T, following experience E. Depending on how we specify T, P, and E, the learning task might also be called by names such as data mining, autonomous discovery, database updating, programming by example, etc.

Machine learning is a computer’s way of using a given data set to figure out how to perform a specific function through trial and error.





What is a specific function? A simple example is deciding the best e-mail subject line for people who used certain search terms to find your website, their behavior on your website, and their subsequent responses (or lack thereof) to your e-mails.

The machine looks at previous results, formulates a conclusion, and then waits for the results of a test of its hypothesis. The machine next consumes those test results and updates its weighting factors from which it suggests alternative subject lines—over and over.

There is no final answer because reality is messy and ever changing. So, just like humans, the machine is always accepting new input to formulate its judgments. It's learning.

The “three *Ds*” of artificial intelligence are that it can *detect*, *decide*, and *develop*.

Detect

AI can discover which elements or attributes in a subject matter domain are the most predictive. Even with a great deal of noisy data and a large variety of data types, it can identify the most revealing characteristics, figuring out which to heed to and which to ignore.

Decide

AI can infer rules about data, from the data, and weigh the most predictive attributes against each other to make a decision. It can take an enormous number of characteristics into consideration, ponder the relevance of each, and reach a conclusion.

Develop

AI can grow and mature with each iteration. Whether it is considering new information or the results of experimentation, it can alter its opinion about the environment as well as how it evaluates that environment. It can program itself.

WHOM IS THIS BOOK FOR?

This is the sort of book data scientists should buy for their marketing colleagues to help them understand what goes on in the data science department.





6 ▶ ARTIFICIAL INTELLIGENCE FOR MARKETING

This is the sort of book marketing professionals should buy for their data scientists to help them understand what goes on in the marketing department.

This book is for the marketing manager who has to respond to the C-level insistence that the marketing department “get with the times” (management by *in-flight* magazine).

This book is for the marketing manager who has finally become comfortable with analytics as a concept, and learned how to become a dexterous consumer of analytics outputs, but must now face a new educational learning curve.

This book is for the rest of us who need to understand the big, broad brushstrokes of this new type of data processing in order to understand where we are headed in business.

This book is for those of us who need to survive even though we are not data scientists, algorithm magicians, or predictive analytics statisticians.

We must get a firm grasp on artificial intelligence because it will be our jobs to make use of it in ways that raise revenue, lower costs, increase customer satisfaction, and improve organizational capabilities.



THE BRIGHT, BRIGHT FUTURE



Artificial intelligence will give you the ability to match information about your product with the information your prospective buyers need at the moment and in a format they are most likely to consume it most effectively.

I came across my first seemingly self-learning computer system when I was selling Apple II computers in a retail store in Santa Barbara in 1980. Since then, I’ve been fascinated by how computers can be useful in life and work. I was so interested, in fact, that I ended up explaining (and selling) computers to companies that had never had one before, and programming tools to software engineers, and consulting to the world’s largest corporations on how to improve their digital relationships with customers through analytics.

Machine learning offers so much power and so much opportunity that we’re in the same place we were with personal computers in 1980, the Internet in 1993, and e-commerce when Amazon.com began taking over e-commerce.

In each case, the promise was enormous and the possibilities were endless. Those who understood the impact could take advantage of it before their competitors. But the advantage was fuzzy, the implications were diverse, and speculations were off the chart.





through, narrow machine intelligence is good enough for problem subsets that are incredibly routine.

We have so many companies that are dedicated to marketing automation and to smart agents and smart bots. If we were to enumerate all the jobs being done in marketing department and score them based on how much pain caused, and how esteemed they are, you'd have no shortage of start-ups trying to provide the next wave of mechanization in the age of information.

And heaven knows, we have plenty of well-paid people spending a great deal of time doing incredibly routine work.

So machine learning is great. It's powerful. It's the future of marketing. But just what the heck *is* it?

WHAT'S ALL THIS AI THEN?

What are AI, cognitive computing, and machine learning? In "The History of Artificial Intelligence,"⁸ Chris Smith introduces AI this way:

The term *artificial intelligence* was first coined by John McCarthy in 1956 when he held the first academic conference on the subject. But the journey to understand if machines can truly think began much before that. In Vannevar Bush's seminal work *As We May Think* (1945) he proposed a system which amplifies people's own knowledge and understanding. Five years later Alan Turing wrote a paper on the notion of machines being able to simulate human beings and the ability to do intelligent things, such as play Chess (1950).

In brief—AI mimics humans, while machine learning is a system that can figure out how to figure out a specific task. According to SAS, multinational developer of analytics software, "Cognitive computing is based on self-learning systems that use machine-learning techniques to perform specific, humanlike tasks in an intelligent way."⁹

THE AI UMBRELLA

We start with *AI*, *artificial intelligence*, as it is the overarching term for a variety of technologies. AI generally refers to making computers act like people. "Weak AI" is that which can do something very specific,





10 ► ARTIFICIAL INTELLIGENCE FOR MARKETING

very well, and “strong AI” is that which thinks like humans, draws on general knowledge, imitates common sense, threatens to become self-aware, and takes over the world.

We have lived with weak AI for a while now. Pandora is very good at choosing what music you might like based on the sort of music you liked before. Amazon is pretty good at guessing that if you bought *this*, you might like to buy *that*. Google’s AlphaGo beat Go world champion Lee Sedol in March 2016. Another AI system (DeepStack) beat experts at no-limit, Texas Hold’em Poker.¹⁰ But none of those systems can do anything else. They are *weak*.

Artificial intelligence is a large umbrella. Under it, you’ll find visual recognition (“That’s a cat!”), voice recognition (you can say things like, “It won’t turn on” or “It won’t connect to the Internet” or “It never arrived”), natural language processing (“I think you said you wanted me to open the garage door and warm up your car. Is that right?”), expert systems (“Based on its behavior, I am 98.3% confident that is a cat”), affective computing (“I see cats make you happy”), and robotics (I’m acting like a cat).

THE MACHINE THAT LEARNS

The magic of machine learning is that it was designed to learn, not to follow strict rules. This is the most fundamental aspect to understand and the most important to remember when you hit that inevitable frustration when things start going slightly off-track. A rules-based system does exactly what it’s told and nothing more. We are comforted by that. A command to send out a gazillion e-mails with the “<first_name>” after the salutation does precisely that. That’s good.

Of course, when the database has something fishy in the first_name field, then somebody gets an e-mail that begins, “Hello, Null, how are you?” Once humans know to look for those sorts of mistakes, we create processes to check and correct the data before hitting Send the next time.

When a batch of e-mails goes out that *all* say, “Hello, <first_name>, how are you?” and the e-mails all include those brackets and that underline, we know to flail the programmers until they find the errant semicolon that caused the problem. In both cases, we can backtrack, find the problem, and fix it.

Machine learning has data and is given a task. It performs the task and waits for feedback. If the feedback is positive, it tries to duplicate the same actions the next time for a similar task. If the feedback is negative, it requires information on where it went wrong.





Guess the Animal

In 1980, at the Personal Electronics store in Santa Barbara, California, I was shown a program that could guess the animal you had in mind. I was flabbergasted. It was a simple case of Asimov's "Any sufficiently advanced technology is indistinguishable from magic."

It: Think of an animal . . .

Is it an elephant? (y or n)

Me: n

It: You win. Help me learn from my mistake before you go . . . What animal were you thinking of?

Me: a rabbit

It: Give me a question to distinguish a rabbit from an elephant.

Me: Is it a small animal?

It: For a rabbit, what is the answer to your question? (y or n)

Me: y

It: Thanks.

Play again? (y or n)

Me: y

It: Think of an animal . . .

Is it a small animal? (y or n)

Me: y

It: Is it a rabbit? (y or n)

Me: n

It: You win. Help me learn from my mistake before you go . . . What animal were you thinking of?

Me: a frog

It: Give me a question to distinguish a frog from rabbit and an elephant.

Me: Is it green?





12 ► ARTIFICIAL INTELLIGENCE FOR MARKETING

It: Thanks.

Play again? (y or n)

Me: No, I see how this works now.

It: I don't understand.

Play again? (y or n)

Me: n

It: Goodbye

Is this program rudimentary? You bet. Was this machine learning? Almost.

After running again and again, the game could guess exactly what animal you had in mind after only a few questions. It was impressive, but it was just following programmed logic. It was not *learning*. Guess the Animal could update its rules-based database and *appear* to be getting smarter as it went along, but it did not change how it made decisions.

The Machine that Programs Itself

Machine learning systems look for patterns and try to make sense of them. It all starts with the question: What problem are you trying to solve?

Let's say you want the machine to recognize a picture of a cat. Feed it all the pictures of cats you can get your hands on and tell it, "These are cats." The machine looks through all of them, looking for patterns. It sees that cats have fur, pointy ears, tails, and so on, and waits for you to ask a question.

"How many paws does a cat have?"

"On average, 3.24."

That's a good, solid answer from a regular database. It looks at all the photos, adds up the paws, and divides by the number of pictures.

But a machine learning system is designed to learn. When you tell the machine that most cats have four paws, it can "realize" that it cannot see all of the paws. So when you ask,

"How many ears does a cat have?"

"No more than two."





the machine has learned something from its experience with paws and can apply that learning to counting ears.

The magic of machine learning is building systems that build themselves. We teach the machine to learn how to learn. We build systems that can write their own algorithms, their own architecture. Rather than learn more information, they are able to change their minds about the data they acquire. They alter the way they perceive. They learn.

The code is unreadable to humans. The machine writes its own code. You can't fix it; you can only try to correct its behavior.

It's troublesome that we cannot backtrack and find out where a machine learning system went off the rails if things come out wrong. That makes us decidedly uncomfortable. It is also likely to be illegal, especially in Europe.

"The EU General Data Protection Regulation (GDPR) is the most important change in data privacy regulation in 20 years" says the homepage of the EU GDPR Portal.¹¹ Article 5, Principles Relating to Personal Data Processing, starts right out with:

Personal Data must be:

- * processed lawfully, fairly, and in a manner transparent to the data subject
- * collected for specified, explicit purposes and only those purposes
- * limited to the minimum amount of personal data necessary for a given situation
- * accurate and where necessary, up to date
- * kept in a form that permits identification of the data subject for only as long as is necessary, with the only exceptions being statistical or scientific research purposes pursuant to article 83a
- * Parliament adds that the data must be processed in a manner allowing the data subject to exercise his/her rights and protects the integrity of the data
- * Council adds that the data must be processed in a manner that ensures the security of the data processed under the responsibility and liability of the data controller

Imagine sitting in a bolted-to-the-floor chair in a small room at a heavily scarred table with a single, bright spotlight overhead and a detective leaning in asking, "So how did your system screw





16 ▶ ARTIFICIAL INTELLIGENCE FOR MARKETING

strategic defense. Skynet begins to learn at a geometric rate. It becomes self-aware at 2:14 a.m. Eastern time, August 29th. In a panic, they try to pull the plug.

The Terminator, Orion Pictures, 1984

At the end of 2014, Professor Stephen Hawking rattled the data science world when he warned, “The development of full artificial intelligence could spell the end of the human race It would take off on its own, and re-design itself at an ever increasing rate. Humans, who are limited by slow biological evolution, couldn’t compete and would be superseded.”¹⁶

In August 2014, Elon Musk took to Twitter to express his misgivings:

“Worth reading Superintelligence by Bostrom. We need to be super careful with AI. Potentially more dangerous than nukes,” (Figure 1.2) and “Hope we’re not just the biological boot loader for digital superintelligence. Unfortunately, that is increasingly probable.”

In a clip from the movie *Lo and Behold*, by German filmmaker Werner Herzog, Musk says:

I think that the biggest risk is not that the AI will develop a will of its own, but rather that it will follow the will of people that establish its utility function. If it is not well thought out—even if its intent is benign—it could have quite a bad outcome. If you were a hedge fund or private equity fund and you said, “Well, all I want my AI to do is



Figure 1.2 Elon Musk expresses his disquiet on Twitter.





maximize the value of my portfolio," then the AI could decide, well, the best way to do that is to short consumer stocks, go long defense stocks, and start a war. That would obviously be quite bad.

While Hawking is thinking big, Musk raises the quintessential Paperclip Maximizer Problem and the Intentional Consequences Problem.

The AI that Ate the Earth

Say you build an AI system with a goal of maximizing the number of paperclips it has. The threat is that it learns how to find paperclips, buy paperclips (requiring it to learn how to make money), and then work out how to manufacture paperclips. It would realize that it needs to be smarter, and so increases its own intelligence in order to make it even smarter, in service of making paperclips.

What is the problem? A hyper-intelligent agent could figure out how to use nanotech and quantum physics to alter all atoms on Earth into paperclips.

Whoops, somebody seems to have forgotten to include the Three Laws of Robotics from Isaac Asimov's 1950 book, *I Robot*:

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 it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

Max Tegmark, president of the Future of Life Institute, ponders what would happen if an AI

is programmed to do something beneficial, but it develops a destructive method for achieving its goal: This can happen whenever we fail to fully align the AI's goals with ours, which is strikingly difficult. If you ask an obedient intelligent car to take you to the airport as fast as possible, it might get you there chased by helicopters and covered in vomit, doing not what you wanted but literally what you asked for. If a superintelligent system is tasked with a(n) ambitious geoengineering project, it might wreak havoc with our ecosystem as a side effect, and view human attempts to stop it as a threat to be met.¹⁷





If you really want to dive into a dark hole of the existential problem that AI represents, take a gander at “The AI Revolution: Our Immortality or Extinction.”¹⁸

Intentional Consequences Problem

Bad guys are the scariest thing about guns, nuclear weapons, hacking, and, yes, AI. Dictators and authoritarian regimes, people with a grudge, and people who are mentally unstable could all use very powerful software to wreak havoc on our self-driving cars, dams, water systems, and air traffic control systems. That would, to repeat Mr. Musk, obviously be quite bad.

That’s why the Future of Life Institute offered “Autonomous Weapons: An Open Letter from AI & Robotics Researchers,” which concludes, “Starting a military AI arms race is a bad idea, and should be prevented by a ban on offensive autonomous weapons beyond meaningful human control.”¹⁹

In his 2015 presentation on “The Long-Term Future of (Artificial) Intelligence,” University of California, Berkeley professor Stuart Russell asked, “What’s so bad about the better AI? AI that is incredibly good at achieving something other than what we *really* want.”

Russell then offered some approaches to managing the it’s-smarter-than-we-are conundrum. He described AIs that are not in control of anything in the world, but only answer a human’s questions, making us wonder whether it could learn to manipulate the human. He suggested creating an agent whose only job is to review other AIs to see if they are potentially dangerous and admitted that was a bit of a paradox. He’s very optimistic, however, given the economic incentive for humans to create AI systems that do *not* run amok and turn people into paperclips. The result will inevitably be the development of community standards and a global regulatory framework.

Setting aside science fiction fears of the unknown and a madman with a suitcase nuke, there are some issues that are real and deserve our attention.

Unintended Consequences

The biggest legitimate concern facing marketing executives when it comes to machine learning and AI is when the machine does what you tell it to do rather than what you wanted it to do. This is much like the paperclip problem, but much more subtle. In broad terms, this





is known as the *alignment problem*. The alignment problem wonders how to explain to an AI system goals that are not absolute, but take all of human values into consideration, especially considering that values vary widely from human to human, even in the same community. And even then, humans, according to Professor Russell, are irrational, inconsistent, and weak-willed.

The good news is that addressing this issue is actively happening at the industrial level. “OpenAI is a non-profit artificial intelligence research company. Our mission is to build safe AI, and ensure AI’s benefits are as widely and evenly distributed as possible.”²⁰

The other good news is that addressing this issue is actively happening at the academic/scientific level. The Future of Humanity Institute teamed with Google to publish a paper titled “Safely Interruptible Agents.”²¹

Reinforcement learning agents interacting with a complex environment like the real world are unlikely to behave optimally all the time. If such an agent is operating in real-time under human supervision, now and then it may be necessary for a human operator to press the big red button to prevent the agent from continuing a harmful sequence of actions—harmful either for the agent or for the environment—and lead the agent into a safer situation. However, if the learning agent expects to receive rewards from this sequence, it may learn in the long run to avoid such interruptions, for example by disabling the red button—which is an undesirable outcome. This paper explores a way to make sure a learning agent will not learn to prevent (or seek!) being interrupted by the environment or a human operator. We provide a formal definition of safe interruptibility and exploit the off-policy learning property to prove that either some agents are already safely interruptible, like Q-learning, or can easily be made so, like Sarsa. We show that even ideal, uncomputable reinforcement learning agents for (deterministic) general computable environments can be made safely interruptible.

There is also the Partnership on Artificial Intelligence to Benefit People and Society,²² which was “established to study and formulate best practices on AI technologies, to advance the public’s understanding of AI, and to serve as an open platform for discussion and engagement about AI and its influences on people and society.”





Granted, one of its main goals from an industrial perspective is to calm the fears of the masses, but it also intends to “support research and recommend best practices in areas including ethics, fairness, and inclusivity; transparency and interoperability; privacy; collaboration between people and AI systems; and of the trustworthiness, reliability, and robustness of the technology.”

The Partnership on AI’s stated tenets²³ include:

We are committed to open research and dialog on the ethical, social, economic, and legal implications of AI.

We will work to maximize the benefits and address the potential challenges of AI technologies, by:

Working to protect the privacy and security of individuals.

Striving to understand and respect the interests of all parties that may be impacted by AI advances.

Working to ensure that AI research and engineering communities remain socially responsible, sensitive, and engaged directly with the potential influences of AI technologies on wider society.

Ensuring that AI research and technology is robust, reliable, trustworthy, and operates within secure constraints.

Opposing development and use of AI technologies that would violate international conventions or human rights, and promoting safeguards and technologies that do no harm.

That’s somewhat comforting, but the blood pressure lowers considerably when we notice that the Partnership includes the American Civil Liberties Union. That makes it a little more socially reliable than the Self-Driving Coalition for Safer Streets, which is made up of Ford, Google, Lyft, Uber, and Volvo without any representation from little old ladies who are just trying to get to the other side.

Will a Robot Take Your Job?

Just as automation and robotics have displaced myriad laborers and word processing has done away with legions of secretaries, some jobs will be going away.

The *Wall Street Journal* article, “The World’s Largest Hedge Fund Is Building an Algorithmic Model from Its Employees’ Brains,”²⁴ reported





Christopher Berry sees a threat to the lower ranks of those in the marketing department.²⁸

If we view it as being a way of liberating people from the drudgery of routine within marketing departments, that would be quite a bit more exciting. People could focus on the things that are most energizing about marketing like the creativity and the messaging—the stuff people enjoy doing.

I just see nothing but opportunity in terms of tasks that could be automated to liberate humans. On the other side, it's a typical employment problem. If we get rid of all the farming jobs, then what are people going to do in the economy? It could be a tremendous era of a lot more displacement in white collar marketing departments.

Some of the first jobs to be automated will be juniors. So we could be very much to a point where the traditional career ladder gets pulled up after us and that the degree of education and professionalism that's required in marketing just increases and increases.

So, yes, if you've been in marketing for a while, you'll keep your job, but it will look very different, very soon.

MACHINE LEARNING'S BIGGEST ROADBLOCK

That would be *data*. Even before the application of machine learning to marketing, the glory of *big data* was that you could sort, sift, slice, and dice through more data than previously computationally possible.

Massive numbers of website interactions, social engagements, and mobile phone swipes could be sucked into an enormous database in the cloud and millions of small computers that are so much better, faster, and cheaper than the Big Iron of the good old mainframe days could process the heck out of it all. The problem then—and the problem now—is that these data sets do not play well together.

The best and the brightest data scientists and analysts are still spending an enormous and unproductive amount of time performing janitorial work. They are ensuring that new data streams are properly vetted, that legacy data streams continue to flow reliably, that the data





24 ▶ ARTIFICIAL INTELLIGENCE FOR MARKETING

that comes in is formatted correctly, and that the data is appropriately groomed so that all the bits line up.

- Data set A starts each week on Monday rather than Sunday.
- Data set B drops leading zeros from numeric fields.
- Data set C uses dashes instead of parentheses in phone numbers.
- Data set D stores dates European style (day, month, year).
- Data set E has no field for a middle initial.
- Data set F stores transaction numbers but not customer IDs.
- Data set G does not include in-page actions, only clicks.
- Data set H stores a smartphone's IMEI or MEID number rather than its phone number.
- Data set I is missing a significant number of values.
- Data set J uses a different scale of measurements.
- Data set K, and so on.

It's easy to see how much work goes into data cleansing and normalization. This seems to be a natural challenge for a machine learning application.

Sure enough, there are academics and data scientists working on this, but they're a long way off. How can you tell?

In their paper titled "Probabilistic Noise Identification and Data Cleaning,"²⁹ Jeremy Kubica and Andrew Moore describe their work on *not* throwing out entire records when only some of the fields are contaminated. "In this paper we present an approach for identifying corrupted fields and using the remaining non-corrupted fields for subsequent modeling and analysis. Our approach learns a probabilistic model from the data that contains three components: a generative model of the clean data points, a generative model of the noise values, and a probabilistic model of the corruption process."

It's a start.

MACHINE LEARNING'S GREATEST ASSET

That would be *data*. Machine learning has a truly tough time with too little information. If you give it only one example, it can tell you *exactly* what to expect the next time with 100 percent confidence. It will be wrong.

Machine learning doesn't work like statistics. Statistics can tell you the likelihood of a coin toss or the probability of a plane crash.





PROBABILITY OF A PLANE CRASH

Three statisticians are in a plane when the pilot announces that they've lost one of their engines. "But it's okay, folks, these planes were built to fly under the worst conditions. It does mean, however, that we're going to fly a bit slower and we'll be about a half an hour late. Please don't worry. Sit back, relax, and enjoy the rest of your flight."

The first statistician says, "There's still a 25 percent chance that I'll make my connection."

Fifteen minutes later, the pilot is on the PA again. "Ladies and gentlemen, we seem to have lost a second engine. No problem, the others are still going strong. This does mean, however, that we'll be about an hour late to the gate. I'm so sorry for the inconvenience."

The second statistician says, "There's an 83 percent chance I'm going to miss my dinner."

After a half an hour, the pilot makes *another* announcement, "Ladies and Gents, we've lost yet another engine. Yes, I know this is bad, but there's really no need to worry. We'll make it just fine, but we're going to be two hours late to the airport."

The third statistician says, "That last engine better not fail or we'll *never* land!"

Human experience and ingenuity have worked wonders for marketing for hundreds of years: gut feel and common sense. When we added statistics to the mix, we expanded our experience by considering historical precedent. But we still rely on gut feel as we feel around blindly in the data, hoping to stumble on something recognizable.

How We Used to Dive into Data

As the Board Chair of the Digital Analytics Association, I strove to explain how digital analysts go beyond answering specific questions. I wrote the following in the Applied Marketing Analytics Journal, describing the role of the "data detective."

Discovering Discovery, Data Discovery Best Practices³⁰

A crystal ball is filled with nothing at all or smoke and clouds, mesmerizing the uninitiated, but very useful for the scrying specialist. The crystal ball mystic is tasked with entertaining more than communicating genuine visions. Creating something from nothing takes imagination, creativity, and the ability to read





one's fellow man to determine what fictions they might consider valuable. The medium who directs a séance is in much the same role.

Tarot Card readers are a step closer to practicality. They use their cards as conversation starters. "You drew The Magician, which stands for creation and individuality, next to the Three of Cups, which represents a group of people working together. Are you working on a project with others right now?" The "mystical" conversation is all about the subject, and therefore, seems revelatory.

The Digital Analyst also has a crystal ball (The Database) and Tarot Cards (Correlations) with which to entice and enthrall the Truth Seeker. The database is a mystery to the supplicant, and the correlations seem almost magical.

The Digital Analyst has something more powerful than visions and more practical than psychology—although both are necessary in this line of work. The analyst has data; data that can be validated and verified. Data that can be reliably used to answer specific questions.

The Digital Analyst truly shines when seeking insight beyond the normal, predictable questions asked on a daily basis. The analyst can engage in discovery; the art of uncovering important truths that can be useful or even transformative to those who would be data-driven.

Traditional Approach: Asking Specific Questions

A business manager wants to know the buying patterns of her customers.

A shipping manager wants to project what increased sales will mean to staffing.

A production manager wants to anticipate and accordingly adjust the supply chain.

An advertising professional wants to see the comparative results of a half a dozen promotional campaigns.

Each of these scenarios call for specific data to be assembled and tabulated to provide a specific answer. Proper data collecting, cleansing, and blending are required, and can be codified if the same questions are to be asked repeatedly. And thus, reporting is born.





Reports are valuable and necessary . . . until they are not. Then they are the source of repetitive stress, adding no value to the organization. The antidote is discovery.

Exploring Data

An investigation is an effort to get data to reveal what it knows. (“Where were you on the night of the 27th?”). But data discovery is the art of interviewing data to learn things you didn’t necessarily know you wanted to know.

The Talented data explorer is much like the crystal ball gazer and the Tarot reader in several ways. They:

- Have a method for figuring out what the paying customer wants to know.
- Have broad enough knowledge about the subject to recognize potentially interesting details.
- Are sufficiently open minded to be receptive to details that *might* be relevant.
- Keep in close communication with the petitioner to guide the conversation.
- Understand the underlying principles well enough to push the boundaries.
- Are curious by nature and enjoys the intellectual hunt.

Data discovery is part mind reading, part pattern recognition, and part puzzle solving. Reading the mind of the inquisitor is obligatory to ensure the results are of interest to those with control of the budget. Pattern recognition is a special skill that can be honed to help direct lines of enquiry and trains of thought. An aptitude for detective work is the most important talent of the Digital Analyst; that ability to ponder the meaning of newly uncovered evidence.

Data discovery is the art of mixing an infinitely large bowl of alphabet soup and being able to recognize the occasional message that floats to the surface in an assortment of languages. Although, with Big Data, adding more data variety to the mix, the Digital Analyst must also be able to read tea leaves, translate the I Ching, generate an astrological chart, interpret dreams, observe auras, speak in tongues, and sing with sirens in order to turn lead into gold.





TDWI Research finds that enterprise data warehouses, BI reporting and OLAP cubes, spreadsheets, and analytic databases are the most important data sources for visual analysis and data discovery, according to survey respondents. (TDWI⁴)

The care and feeding of the raw material used in the data discovery process is even more important in light of the lack of five-star chefs. As analytics becomes more accepted, demanded and democratized, more and more amateur analysts will be deriving conclusions from raw material they trust implicitly rather than understand thoroughly. Preparing for data illiterate explorers requires even more rigor than usual to guard against their impulse to jump to the wrong conclusions.

Asking Really Good Questions

In the hands of a well-informed analyst, lots of data and heavy-lifting analytics tools are very powerful. Getting the most out of this combination takes a little bit of creativity.

Creativity means broadening your mental scope. Rather than seeking a specific answer, open yourself up to possibilities. It's like focusing on your peripheral vision.

1. Appreciate Anomalies

Whether you use visualization tools and “look for” things that go bump in the night, or you are adept at scanning a sea of numbers and wondering why it looks out of balance, the skill to hone is the art of seeing the out-of-the-ordinary.

Outliers, spikes, troughs—any anomaly—are our friends. They draw our attention to that which is not like the others and spark the intellectual exercise of wondering “Why?”

What is it about this element that makes it point in a different direction? Could it be some error in the collection or transformation of the underlying data? Is it a function of how the report was written or the query was structured? Or does it represent some new behavior/market movement/customer trend?

It is in the hunt for the truth about these standouts that we trip over the serendipitous component that spawns a new





question and another dive down the rabbit hole.
The secret is knowing when to stop.

One can easily get lost in a hyperlink-chasing “research session” and burn hours with very little to show for it. Following the scent of significance is an art and one that takes practice and discipline. Many scientists spend a career pursuing a specific outcome only to find it disproved. Others stop just short of a discovery because they lose heart. The magic happens between those two points.

Give in to the temptation to slice the data one more time or to cross reference results against just one more query, but be vigilant that you are not wasting valuable cycles on diminishing returns.

If you don’t see what you expect to see, work your hardest to understand why. It may be that you do not have enough facts. It might be that you have already, unknowingly, come to a conclusion or formed a pet theory without all the facts. It might be—and this is the most likely—that there is something afoot which you have not yet considered.

Dig deeper. Ask, “I wonder . . .” And be cognizant of that which is conspicuous in its absence.

Gregory (Scotland Yard detective): “Is there any other point to which you would wish to draw my attention?”

Holmes: “To the curious incident of the dog in the night-time.”

Gregory: “The dog did nothing in the night-time.”

Holmes: “That was the curious incident.”

Sir Arthur Conan Doyle, Silver Blaze

As a corollary, be wary of the homologous as well:

1. Exhibiting a degree of correspondence or similarity.
2. Corresponding in structure and evolutionary origin, but not necessarily in function.

For example, human arm, dog foreleg, bird wing, and whale flipper are homologous. (A Word A Day⁵)





Things that are unusually similar are equally cause for alarm as standouts. If everybody in your cohort looks the same, there's something funny going on and it's worth an investigation. It may be that their similarity is a statistical anomaly.

2. Savor Segmentation

People (thank heaven!) are different. We make a huge mistake when we lump them all together. But we cannot treat them as individuals—yet. Peppers and Rogers' *One to One Future* is not yet upon us. In between lies segmentation.

It almost doesn't matter how you segment your customers (geographically, chronologically, by hair color). Eventually, you will find traits that are useful in finding a cluster of behavior that can be leveraged to your advantage.

People who come to our website in the morning are more likely to X.

People who complain about us on social media respond better to message Y.

People who use our mobile app more than twice a week are more likely to Z.

When it comes to segmenting customers by behavior, Bernard Berelson pretty much nailed it in his "Human Behavior: An Inventory of Scientific Findings"⁶ where he said:

Some do and some don't.

The differences aren't that great.

It's more complicated than that.

When you're trying to get the right message in front of the right people at the right time and on the right device, segmentation may likely be the key to the mystery.

3. Don't Fool Yourself

While working with data is reassuring—we are, after all dealing with facts and not opinions—we are still human and still faced with serious mental handicaps.





Being open-minded and objective are wonderful goals, but they are not absolute.

Cognitive biases are inherited, taught, and picked up by osmosis in a given culture. In short, your mind can play tricks on you. While this is too large a subject to cover in depth here, there are some examples that make it clear just how tenuous your relationship with “the facts” might be.

Familiarity Bias

I’ve worked in television advertising all my life and I can tell you without any doubt that it’s the most powerful branding medium there is.

Hindsight or Outcome Bias

If they’d only have asked me, I would have told them that the blue button would not convert as well as the red one. It was obvious all along.

Attribution Bias

Of course I should have turned left at that light. But I was distracted by the sun in my eyes and the phone ringing. That other guy missed the turn because he’s a dim-wit.

Representativeness Bias

Everybody who clicks on that link must be like everybody else who clicked on that link in the past.

Anchoring Bias

That’s far too much to pay for this item. The one next to it is half the price.

Availability Bias (the first example that comes to mind)

That’ll never work—let me tell you what happened to my brother-in-law . . .

Bandwagon Bias

We should run a Snapchat campaign because everybody else is doing it.

Confirmation Bias

I’m a conservative, so I only watch Fox News.

I’m a liberal, so I only watch The Rachel Maddow Show.

I’ve been in advertising all my life, so I count on Nielsen, Hitwise, and comScore.





34 ► ARTIFICIAL INTELLIGENCE FOR MARKETING

I started out grepping log files, so I only trust my Coremetrics/Omniture/Webtrends numbers.

Projection Bias

I would never click on a product demo without a long list of testimonials, so we can assume that's true of everybody else.

Expectancy Bias

Your report must be wrong because it does not show the results I anticipated.

Normalcy Bias

Back-ups? We've never had a data loss problem yet, I don't see it happening this quarter so we won't have to budget for it.

Semmelweis Reflex

I don't care what your numbers say, we've always had better conversions from search than social media so we're not going to change our investment.

If any of the above sound familiar, congratulations—you've been paying attention. The hard part is convincing others that there may be a cognitive problem.

4. Correlation versus Causation

While frequently mentioned, it cannot be stressed enough that just because drownings go up when ice cream sales go up, one did not cause the other.

Most recently, a Swedish study ("Allergy in Children in Hand Versus Machine Dishwashing"⁷) concluded, "In families who use hand dishwashing, allergic diseases in children are less common than in children from families who use machine dishwashing," and speculated that, "a less-efficient dishwashing method may induce tolerance via increased microbial exposure."

While the study asked a great deal of questions about the types of food they eat, food preparation, parental smoking, etc., there are simply too many other variables at play for this cause to be solely responsible for that effect. How many other similarities are there among families that have dishwashers vs. those that do not?



6. Become a Change Agent

The very best way to win the hearts and minds of those who can benefit the most by your flair for data discovery is to educate them.

The more people in your organization who understand the ways and means of data exploration as well as the associate risks and rewards, the more they will come to you for answers, include you in planning sessions and support your calls for more data, people and tools.

Start by inviting them to lunch. Ask them to bring their best questions about The Data. Encourage those who would rather not be seen as ill-informed to submit their questions in advance. Prepare a handful of questions that you wish they would ask.

Answer their questions. Show them examples of quick-wins enjoyed by other projects in other departments. Share case studies from vendors about successes at other companies.

Engage your audience in the excitement of the chase with a simple data set and a common challenge. If you can teach them how to ask great questions by example and by exercise then you can change how they approach data—to see it as a tool instead of an accusation.

And be sure to feed them. This is a case where a free lunch will pay off handsomely.

Your Job as Translator

You know your data inside and out, but the consumers of your insights, who must depend on your recommendations do not. To them, your data is as readable as a crystal ball or a sequence of Tarot cards. That means they are putting their trust in you.

Therefore, your responsibility is to inform without confusing, to encourage without mystifying and to reassure without resorting to sleight of hand. Entice and enthrall your Truth Seekers with The Data and The Correlations, but make sure your confidence levels are high and be prepared to show your work.





Conclusion

Successful data discovery requires good tools (technology) and trustworthy raw material (clean data), but depends on the creativity of the data detective. The best analyst has the ability to manipulate data in a variety of ways to tease out relevant insights. With the goals of the organization firmly in mind, top analysts engage the data in a conversation of What-Ifs, resulting in tangible insights that can be used to make decisions by those in charge. The analyst, as consulting detective, becomes indispensable.

NOTES

1. Self-Service Data Discovery and Visualization Application, Sense BI Tool | Qlik, available at <http://www.qlik.com/us/explore/products/sense>, last accessed on 3/13/15.
2. Data Discovery | Tableau Software, available at <http://www.tableau.com/solutions/data-discovery>, last accessed on 3/13/15.
3. Features of the Analytics Platform | MicroStrategy, available at <http://www.microstrategy.com/us/analytics/features>, last accessed on 3/13/15.
4. Data Visualization and Discovery for Better Business Decisions, available at http://www.adaptiveinsights.com/uploads/news/id421/tdwi_data_visualization_discovery_better_business_decisions_adaptive_insights.pdf, last accessed on 3/13/15.
5. A.Word.A.Day—homologous, available at <http://wordsmith.org/words/homologous.html>, last accessed 3/13/15.
6. Human Behavior: An Inventory of Scientific Findings, available at <http://home.uchicago.edu/aabbott/barbpapers/barbhuman.pdf>, last accessed 3/13/15.
7. Allergy in Children in Hand Versus Machine Dishwashing, available at <http://pediatrics.aappublications.org/content/early/2015/02/17/peds.2014-2968.full.pdf>, last accessed 3/13/15.

Variety of Data Is the Spice of Life

Machine learning differs from data diving. It is like putting tens of thousands of statisticians in a black box and throwing in a question. They

