

LEARNING MADE EASY

Dataiku Special Edition

Marketing Artificial Intelligence (AI)

for
dummies[®]
A Wiley Brand



Discover what AI can
achieve for marketing

Learn the most successful
ways to implement AI

Leave behind old
heuristic methods

Brought to
you by



**data
iku**

Luca Massaron

About Dataiku

Dataiku believes that to succeed in the world's rapidly evolving ecosystem, marketing teams — no matter what their industry or size — must use data to continuously innovate and disrupt the business. But radical disruption doesn't come from technology alone; it also comes from uniting people around technology to bring change. This means introducing transparency around data and data processes to empower everyone throughout the enterprise and the marketing team - not just a siloed group — to draw and use insights from data on their own.

Out of this vision, Dataiku was born as the centralized data platform that moves businesses along their path from simple analytics to marketing AI for the modern enterprise. Dataiku puts data in the hands of marketing teams worldwide, ensuring that they play a role in driving and pioneering the AI revolution within the organization.

For more information, visit **www.dataiku.com**.



Marketing Artificial Intelligence (AI)

Dataiku Special Edition

by Luca Massaron

**for
dummies[®]**
A Wiley Brand

Marketing Artificial Intelligence (AI) For Dummies®, Dataiku Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2019 by John Wiley & Sons, Inc., Hoboken, New Jersey

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 Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. 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 <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Dataiku and the Dataiku logo are trademarks or registered trademarks of Dataiku. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-119-59706-3 (pbk); ISBN 978-1-119-59708-7 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Acquisitions Editor:

Katie Mohr

Editorial Manager:

Rev Mengle

Business Development

Representative: Susan Blessing

Production Editor:

Tamilmani Varadharaj

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions	1
Icons Used in This Book.....	2
Beyond the Book.....	2
CHAPTER 1: Keeping Up with a World in Transition	3
Grasping a Shifting Marketing Paradigm	3
Realizing the Power of Data	4
Understanding What AI Can Do for Marketing.....	5
CHAPTER 2: Introducing the Key Concepts of AI	7
Distinguishing the Role of Machine Learning	7
Getting Into Deep Learning for Marketing	9
Defining Artificial Intelligence	9
CHAPTER 3: Choosing the Right Use Case for You	11
Understanding the Use Cases for AI in Marketing.....	11
Tracking and Measuring Marketing Effectiveness.....	12
Contrasting Churn and Increasing Loyalty	13
Suggesting Customers by Recommendations.....	14
Reaching More and More Objectives.....	15
CHAPTER 4: Getting Started with a Great AI Application and Avoiding Pitfalls	17
Setting Off with a Suitable Team	17
Framing the Right Problem and Audience	19
Keeping Up Trust and Interpretability	19
Preparing the Data Properly	20
Getting Into a Virtuous Circle.....	20
Testing Everything for Production.....	21
Adjusting to Concept Drift.....	22

CHAPTER 5: **Learning from Leading Marketing Teams**..... 23

 Listening To Successful Stories..... 23

 Understanding the Secret Ingredients 25

CHAPTER 6: **Ten Things You Need to Know
About Marketing AI** 27

Introduction

In recent years, machine learning (ML) and advanced analytics (under the same umbrella term), data science, and ultimately artificial intelligence (AI), have changed the business landscape in a radical way. The revolution comes from digital devices that pervade our daily lives and produce copious data about everything at a speed never experienced before. Data science translates the information riches into a competitive advantage for companies that are ready to take advantage of it. Companies with smart marketers leverage these resources to increase productivity, efficiency, and profits.

About This Book

This book helps you understand the tools of the trade and the key trends needed for you to start your journey into ML and AI. This book consists of six chapters that explore the following:

- » The revolutionizing power of data and AI (Chapter 1)
- » Key terms of AI and machine learning (Chapter 2)
- » Most interesting marketing use cases in AI (Chapter 3)
- » Suggestions for starting great AI projects (Chapter 4)
- » Successful stories from leading marketing teams (Chapter 5)
- » Ten suggestions about applied AI for marketing (Chapter 6)

Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but this book assumes a few things nonetheless. Mainly, I assume that you're a marketing professional. Ideally, I expect you to be a marketing leader (like the CMO or head of marketing), but the book also considers the needs of individual contributors on a marketing team, including those in marketing or business analytics.

I also assume that you are familiar with the basics of marketing analytics and associated concepts — for example, the ideas of customer churn, customer segmentation, and marketing attribution, just to name a few. Thus, the book won't go into much detail on the basics of these concepts, and it simply concentrates on everything about AI and machine learning that is relevant for marketing.

Icons Used in This Book

Throughout this book, you occasionally see special icons to call your attention to important information. Here's what to expect.



REMEMBER

This icon points out information you should commit to your nonvolatile memory, your gray matter, or your noggin!



TIP

Tips are appreciated, never expected — and I sure hope you appreciate these useful nuggets of information.



WARNING

These alerts point out the stuff your mother warned you about (well, probably not), but they do offer practical advice.

Beyond the Book

There's only so much that can be covered in 32 short pages, so if you find yourself at the end of this book wondering where you can learn more, you can try other books in the Dummies series, such as *Machine Learning For Dummies* or *Artificial Intelligence For Dummies*, both by John Paul Mueller and Luca Massaron.

IN THIS CHAPTER

- » Realizing how the world is changing because of data
- » Figuring out what digital companies like Google achieve with data
- » Seeing a revolution approaching in marketing

Chapter 1

Keeping Up with a World in Transition

This chapter introduces you to the significant changes affecting businesses and consumers due to the widespread adoption of digital devices. This *digital revolution* represents a cultural shift in how people live their lives and interact with others. Because of this shift, large amounts of data are being created, as well as new tools such as machine learning and AI to extract knowledge about customers.

Grasping a Shifting Marketing Paradigm

The last ten years have seen a data revolution. Year by year, the number of electronic devices that record, store, and transmit data over the Internet has increased, and the amount of data available today for analysis for every business is vast.

However, for a business to take advantage of this revolution, it's not enough to own the data and make it available for use. Top businesses get ahead in this new data-driven economy by handling and transforming this growing mass of data to adapt to changing market conditions. Some of them have even moved into automation, which allows AI and machine learning tools to automatically

extract valuable information from data and immediately act on it to turn that information into revenue or competitive advantage in operations, finance, and marketing. Unfortunately, relying on automation isn't the reality for the average business because data ends up siloed in organizations with no one able to use it efficiently.



REMEMBER

The opportunities for implementing data-driven marketing have grown incredibly. However, even though data is widely available, it's often scattered in *silos* (private databases or other sources that lack easy access), and you may lack the tools to retrieve and integrate the information you need.

Companies like Google, Apple, Facebook, and Amazon are pioneers in Western markets, while Baidu, Alibaba, Tencent, and Xiaomi are pioneers in Asian markets. In fact, you can find innovative startups everywhere gaining market share across all industries. These companies were born from business models based on data, and they actually drive their business on advanced analytics. These companies don't simply have a data team to help data-driven decision-making; machine learning permeates all their processes and products. Google, for instance, tried to improve the way they cool down the computers at their data centers by employing the techniques found in top-most companies in the world. They then developed an AI that reduces the energy bill by 40 percent.

To compete and stay relevant in this rapidly changing world, businesses need to embrace the idea of using data and AI for their marketing activities. This means boldly driving business decisions through analysis and algorithms, ensuring data initiatives appear at the top of the agenda, and setting up tools to allow data projects to go from concept to delivery in the fastest time possible. Businesses that ignore this need risk being left behind and becoming irrelevant — taken out of their market share by other companies that have found the way to understand customer needs using these new technologies properly.

Realizing the Power of Data

AI technology has the potential to revolutionize marketing in most businesses. AI isn't just another shiny new toy. AI solutions are more effective than the heuristics of the past in creating

solutions, and they're the only tool able to leverage the incredible amount of data now available in a smart way. You could use heuristic tools of the past for large volumes of data as well, but only AI (and machine learning) solutions are able to generate insights and predictions that are more effective because they adapt according to data. The more data AI and machine learning see, the better they can become, whereas heuristics stay static and never improve. Your business requires the level of automation provided by advanced analytics, machine learning, and AI because neither heuristics nor humans analyze and interact in a smart manner with data at this scale, especially for marketing purposes.

You can find good examples of such data power in companies like Google, Netflix, Uber, and Amazon that are built on a foundation of data. These companies let data and data applications permeate their organization. For instance, Amazon developed systems that recommend products to customers so well that one third of its revenues are due to its recommender system. Similarly, Netflix recommends video and audio to customers based on viewing patterns. These recommendations influence three fourths of the choices made on its platform.

Amazon and Netflix don't rely on sophisticated market research or need complicated marketing plans. Both use the fact that they can talk directly to their customers through Web and mobile apps to provide customized recommendations. They rely on suitable algorithms that understand the customer's buying history, as well as the customer's present behavior, to estimate each customer's needs.



REMEMBER

Some companies originally built around data and its applications already make AI powered by data a reality. However, these companies are the exception, rather than the rule. For most companies, the journey to leveraging AI isn't as natural and has to start somewhere to keep them competitive and profitable.

Understanding What AI Can Do for Marketing

AI's contribution to marketing isn't limited to solutions such as recommender systems that help customers make purchases. Other successful use cases have demonstrated that AI is profitable

and has an impact on the market when organizations use these solutions:

- » Recommender systems help both customers and your sales force by improving interaction with your customers.
- » Real-time forecasting improves sales and distribution.
- » Identifying customers likely to leave your company and addressing their needs to keep their business.
- » Improving SEO results by determining why one page ranks high on a search engine result page and another appears at the bottom.
- » Natural language processing (NLP) to understand customer communication and immediately provide assistance, based on what customers write through e-mail or on social networks.
- » Hyper-targeted advertising that delivers relevant targets based on real-time customer understanding. This option also supports price optimization and lead scoring.
- » Exact marketing attribution that optimizes the way in which you communicate with customers.

After evaluating the list, you will realize that in terms of activities and objectives there is nothing new under the sun, since as a marketer you have already dealt with all these activities before. What is new about AI marketing is the accuracy you have when delivering and executing toward a specific target, the complete (or almost complete) automation of procedures and actions, and the high economic efficiency resulting from using AI in these marketing activities.



TIP

Customers also take advantage of AI progress to automate part of their decision-making, by using consumer decision tools such as price and product comparison websites. This is yet another reason for marketing teams to leverage AI: If you don't, you risk falling out of sync with customers and losing the competitive edge as consumers turn to AI solutions as well.

- » Distinguishing different kinds of artificial intelligence
- » Solving problems using machine learning and data
- » Imagining what deep learning can do for marketing

Chapter 2

Introducing the Key Concepts of AI

This chapter acquaints you with key AI and machine learning terms. In addition to understanding differences behind buzzwords such as AI, machine learning, and deep learning, you also get a grasp of techniques such as supervised and unsupervised learning.

Distinguishing the Role of Machine Learning

Machine learning (ML) isn't new. Some industries have used it since the '90s (CMOS). Many of the companies that are known for their use of ML (such as Netflix, Amazon, and Google) started using it for a wide range of applications in the early 2000s. Even though traditional companies have moved slowly to ML and AI, and ML hasn't really fully broken into these organizations yet, the use of ML to meet specific needs has generally popularized the term and made organizations aware of the software tools that make ML possible.

ML became part of the mainstream business analytics discourse by promising to increase revenues, as it did in companies that

openly state practicing it, such as Google and Amazon. ML turned into a continuously heard buzzword overnight, but you can't define it without some technical insights. From a simplistic view, ML is nothing more than a series of *algorithms* (which are sets of steps that a computer performs to solve a task) that allow computers to learn to solve problems successfully without being specifically programmed to perform the task.

Standard algorithms are strictly devised to complete a specific task (they can't adapt to new problems unless a developer rewrites them). ML algorithms can adapt to a series of different problems because they use data to determine how to tune themselves to the task, without all the human intervention that standard algorithms require.

It's important to understand that not all data is suitable for tuning an ML algorithm. For ML algorithms to work, data should represent the problem and its solution. For instance, you may need to provide suitable images, or detailed customer data and your past business with them. ML then uses algorithms that map between the data you have available and the answer you need.

ML algorithms achieve learning objectives by processing data representing the problem and the answer. Data representing the problem is composed of many different pieces of information called *features*. Features are data representing anything relevant to a problem, such as counting the number of certain shapes in a photo or the amount spent by a customer during an initial transaction with your business. You call data representing your answers *labels*.



TIP

Whether ML can successfully solve problems depends on the features that appear in your data. How you deal with the labels also makes a difference. You need clearly stated labels and historical records that represent them.

Currently, ML can perform two kinds of activities:

» **Supervised learning:** ML transforms data into answers that could be numeric (for instance, the value of future revenues from a customer) or a class (such as, whether a customer will leave you for a competitor), depending on what you need. Getting numbers as a response is a *regression*. When you obtain a class, it's a *classification*.

- » **Unsupervised learning:** You provide ML data without answer labels. In return, you obtain segmentations, similar to the kind used in marketing.

Getting Into Deep Learning for Marketing

Neural networks are state of the art tools for ML. They mimic the human brain (similar to the thinking humanly paradigm in AI) using a network of interconnected mathematical units that automatically crunch and process data. *Deep learning* is just a set of complex neural networks, which emerged in recent years and can achieve more than the neural networks of the past. It's because of deep learning and its *deep neural networks* (DNNs) that it's now possible to:

- » Detect and localize objects in images
- » Transform images into text descriptions
- » Process audio to extract text from it
- » Process text and connect it to a specific meaning
- » Translate foreign languages
- » Interact with humans in an intelligible and reasonable way

All sorts of marketing activities — including distinguishing an image (or face and its expressed feelings), understanding a textual or voice comment from a customer, automatically generating custom content, or providing fine-tuned personalization on websites — are now possible because of deep learning. In the past, similar activities required you to hire a host of experts to study the problem until they could design a piece of computer code that they claimed solved the problem, but not always with satisfying results. Some problems are so complex that you won't get good results no matter how many experts you hire, unless you adopt deep learning as part of your solution.

Defining Artificial Intelligence

It's possible to go on for hours discussing the different definitions of human intelligence. The same goes for defining AI: There are lots of totally different definitions for the term and no one really

agrees or knows what it truly means. Ideally, you want machines to demonstrate intelligence in their behavior (how they act) and in their way of reasoning (why they act in such a way), and you want them to be able to mimic human behavior or even to surpass humans by using higher rationality.



REMEMBER

Misinterpretation of these goals causes many scary predictions, such as when Elon Musk says when using AI we are summoning a demon with superhuman powers.

More practically for marketing purposes, it's possible to say (though that leaves some space for interpretation) that an AI will mostly mimic human skills in perception, learning, interaction, and decision-making:

- » **Perception:** Implies an AI application takes images as an input. Here, ML has an important role because of deep learning and its deep neural networks specialized in image recognition.
- » **Learning:** Specifies the application learns autonomously and effectively from human intervention with respect to the target problem, which also requires ML.
- » **Interaction:** Means that your application can interact with humans in a conversational manner.
- » **Decision-making:** Translates into having an AI application generate advanced and complex strategies to achieve its purposes.

You may need to refer to each of these four skills to find a solution to your marketing problems. The mix you use depends on what you need. For instance, you may need a computer to interact with your customers in a humanlike manner, so the emphasis is on the interaction. As an alternative, you might need to achieve a particular result even if you don't necessarily understand how the computer achieves the goal, which means focusing on decision-making above all (think Google's AlphaGo, an AI that can beat human champions at the ancient Chinese board game using unpredictable moves).



TIP

You can measure AI effectiveness based on the four skills using Florian Douetteau's (who is the CEO at Dataiku) framework for understanding AI. To find out more, go to <https://medium.com/ai-musings/is-it-ai-or-not-a-score-card-with-4-dimensions-8a5297372a55>.

- » Solving the riddle of marketing attribution
- » Keeping your customers loyal
- » Proposing the right purchase suggestions
- » Discovering many other possible use cases for marketing AI

Chapter 3

Choosing the Right Use Case for You

In this chapter, you see the most interesting use cases for AI in marketing, focusing on key applications such as marketing attribution, churn analysis, and recommender systems. You also explore other possible use cases, some already widely used, some up-and-coming.

Understanding the Use Cases for AI in Marketing

Machine learning predicts and segments, while AI automates — leveraging the capabilities of both *algorithms* (able to solve specific problems) and *machine learning* (able to learn how to solve problems that algorithms can't tackle). Marketing already takes advantage of both machine learning and AI to solve ancient marketing problems in a new and radical way.

Tracking and Measuring Marketing Effectiveness

John Wannamaker, a marketing pioneer who lived in the past century, used to say that half the money he invested on advertising was wasted, but the real trouble was he didn't know which half. In the digital age, this isn't a problem anymore. You can track everything and analyze it using an advanced analytical modeling called *marketing attribution*, which can effectively unveil what works and what doesn't in a marketing campaign. Marketing attribution works by measuring the effectiveness of different strategies as part of a digital-only campaign across one specific device (for instance, what sort of ad resulted in a customer purchase).

Attribution isn't new for marketing teams that have dealt with it in the past using heuristic models (such as the *last click heuristic*, where you attribute all the conversion merits to the last media contact or channel you exposed a customer to). Yet, ML-based attribution (and thus AI) is more accurate than the old methods when dealing with the digital world and its vast number of channels. Just think of Facebook, Twitter, Google Ads, and all the other media your customer can interact with. The complexity of the environment illustrates the ineffectiveness of the last click heuristic. ML algorithms have the capability to interrelate all the times a consumer has been exposed to your marketing and determine the exact contribution of each channel in bringing the customer to conversion.



WARNING

Attribution models have some possible pitfalls that you must avoid in order to succeed. Such pitfalls are mostly due to bad planning, and technology can do little to solve them:

- » Not tracking all the channels
- » Improperly tracking across digital devices
- » Forgetting to track the passage from online to offline conversions and vice versa



TIP

The problem is the lack of measurements in key points of consumers' interaction with the campaign: There are blind spots, especially when users use multiple devices, or they show up at the store instead of buying online. These issues could completely compromise the reliability of your marketing attribution analysis.

To avoid such problems, carefully plan and prepare for the customer journey and all the necessary touchpoints in advance. For instance, both digital and physical coupons are a good way to persuade consumers to allow tracking in any online/offline passage they take in their conversion journey.

Contrasting Churn and Increasing Loyalty

When a customer leaves, you aren't simply losing a source of revenue, but also part of your investment, because you invested in a marketing mix to obtain that customer in the first place. Churn analysis (also known as *attrition modeling*) exists to help you prevent profitable customers from leaving without getting a chance to retain them.

Churn analysis relies on careful preparation of customer data, examining available customer history, and taking snapshots at various points in time of the experience. You compare the different snapshots to get an idea whether something has changed during their time as customers. Finally, a machine learning algorithm associates the notable changes to customer actions, such as leaving or staying in your customer base. Used with a lifetime value model, you can power an AI that promptly and automatically reacts when the risk of losing important customers appears. Here are a few pitfalls you have to pay attention to if you're delivering a churn project:

- » Data leakage
- » Targeting the wrong people
- » Thinking of churn analysis as a one-time-project

Similar pitfalls appear in quite a few churn analysis projects. Good planning, and accurate and replicable data preparation, makes it possible to avoid them. *Data leakage* occurs when information not part of the training dataset is used to create the model. Because in churn analysis you work with customer snapshots from the past, the trouble happens when some information you get isn't really from the past. Usually, data leakage happens when your data isn't properly documented and extracted, and you get features created when a customer leaves your company. Detection of data leakage problems isn't easy because data scientists will validate the churn

model as correct, but the model will underperform when put into production.



WARNING

Sometimes results that are too good to be true aren't true at all. Beware of features that are too relevant and critical in your model; usually they're a leakage. Model interpretability tools can reveal the problem. The next chapter of the book tells you more about such tools.

Another common pitfall is to produce models that predict people that are obviously going to churn. This is not only a problem with the model (in particular, with its *recall capability*, which is ability to catch all kinds of churners), but also with the business implications: The output suggests acting to retain people who aren't sensitive to marketing actions because they have decided to leave. Instead, your model should catch the people who are evaluating leaving (the model should act as an early warning system).

Finally, the last pitfall you may see is to perceive the model as a one-shot study. You won't do much with a single iteration of a churn model. Multiple iterations help you create a successful strategy to retain and increase customer loyalty. Moreover, unless you already planned to make your churn analysis ongoing, you'll overlook critical aspects in the project, such as *scalability* (applying it to grow larger customer bases) and *reproducibility* (replicating the same approach in future).

Suggesting Customers by Recommendations

Recommender systems have greatly contributed to the fortunes of businesses such as Amazon and Netflix, to name two. There are different types of recommenders: some based on the historical customer data of a customer you want to advise and some based on the history of all customers. Some recommenders rely on simple processes, while others require demanding math. However, all recommender systems are embedded in some customer process, such as a website that sells goods, as they directly interact with your customers by pointing out products you want the customer to discover. Meanwhile, the recommender system automatically learns to understand a customer's needs better.



Recommender systems are the digital heirs of marketing research. Understanding the needs of customers is the key activity for both approaches, but recommender systems don't use marketing questionnaires. Instead, recommenders use historical data together with customer profiles, and they immediately satisfy customer needs through targeted suggestion.

In spite of their extreme efficacy, recommender systems present some caveats. You must take particular care with:

- » Scalability issues
- » Obvious predictions
- » The cold start problem

Your recommender system should serve all your customers. In addition, it should serve them well — a recommender system that recommends the obvious isn't what your customer needs. As with attribution models, recommender systems are also difficult to test because you can only be completely sure they work or not as they run in a production environment, and people provide satisfaction feedback in an implicit way (they buy or use what you recommend) or explicit way (they rate the suggestions your system provides). Finally, your recommender system may work fine with your present customers, but it may completely fail with new ones, preventing them from becoming better customers, since they don't receive good suggestions. This is the *cold start* problem, when you have to do something with customers you know little about.



There are a few ways to handle a cold start problem. The easiest way, for instance, is recommending the most popular items in stock, as you take time to better understand what you should recommend instead.

Reaching More and More Objectives

So far, this book has focused on AI application contributions to marketing: marketing attribution, attrition modeling, and recommender systems, but there are other relevant options, with even more AI technologies being developed every day. Most of the

applications are already here, ready for you to employ in the right use case in your business:

- » **Hyper-targeted advertising:** Allows you to deliver more relevant ads than ever before through the effective combination and aggregation of new and old data sources.
- » **Pricing optimization:** Helps you dynamically set and update pricing based on a huge variety of ever-changing factors (rather than a few static data points). It makes meeting offer and demand possible at a microlevel.
- » **Lead scoring:** Offers more accurate lead scoring using predictive analytics, which allows you to focus your efforts on prospects that are most likely to buy.
- » **Natural language processing (NLP):** Consists of a series of procedures and techniques that make it easier to process words and phrases to use for statistical analysis, machine learning algorithms, and deep learning. Deep learning especially makes it possible to create conversational AI agents that effectively interact with your customer base or create context for communicating with them.
- » **Real-time forecasting:** Relies on algorithms that take massive data from everywhere and automatically creates reliable sales forecasting. An example of this technology is Facebook's Prophet (<https://facebook.github.io/prophet/>).

- » Choosing the right team and data platform
- » Defining the problem and handling the data
- » Emphasizing trust and interpretability
- » Going into production and avoiding concept drift

Chapter 4

Getting Started with a Great AI Application and Avoiding Pitfalls

This chapter presents a short list of key activities necessary to succeed in creating an effective marketing AI application. It's exciting to skim through the list of actions, but achieving your goals can come with pitfalls that render machine learning and AI ineffective and disillusion many managers. Such trouble may also lie ahead in your implementation.

Setting Off with a Suitable Team

Data science, the umbrella term that includes all data analytics and machine learning (ML) relevant to creating a data product, is a team sport. The team nature of data science comes from the need to create a team composed of specialists with expertise in analytics (they could be experts in data, ML, deep learning, AI, or algorithms) and business experts at the same time. Usually, a data scientist will have technical skills needed to create an AI application, but won't have comparable business knowledge.

Collaboration between data scientists and people intimately familiar with the marketing and business side, even if the latter don't have data skills, is critical for the success of your efforts.

Once you choose members of the team, you must provide them with the right tools to work with. Here you encounter many possible solutions and tools. Don't get into nonsense discussions, such as the benefits of using R or Python as your programming language, or deciding on the best deep learning framework. Keep in mind that tools exist to serve, and each tool serves for a purpose. Reflect instead that you aren't setting up pure R&D, but an applied research team, which means that you have to:

- » Be able to access all the data you need, combining siloed information of different kinds and quality
- » Allow everyone on the team to contribute to the AI implementation through code or analytical activities in a coordinated way
- » Coordinate contributions between team and business people
- » Prepare the data, both features and labels, with the least possible effort and in a fully repeatable way
- » Leverage state-of-the-art algorithms without constraints
- » Painlessly make your solutions usable by the whole organization



TIP

You can achieve these goals only by investing in a data science platform (DS platform), a software tool that orchestrates all parts of the data-to-value process, including cleaning, preparing, and wrangling. Such a system will also oversee the application of ML models, deployment to production, and monitoring of models in production. All of these features should be accessible for coders and noncoders alike and, of course, governable with sound, responsible AI practices.

DS platforms are tools that allow data teams to access and transform data into ML solutions in a cooperative, efficient, replicable, transparent, scalable, and distributable way.



TIP

Setting a proper investment for acquiring a DS platform should be the second activity to plan after deciding on a team. The cost of not having one will soon prove more than the initial investment of getting one.

Framing the Right Problem and Audience

As with more conventional projects you might have contributed to in the past, remember that the more you plan the better. It's essential to put your plan in writing. It's not a waste of time to ask what objectives you want to achieve and what impact you expect to make with your AI marketing application. For instance, when planning to create a churn or an upselling model, it's not academic to check whether a shared definition of what churn or upselling means in your organization exists. Also understanding the audience to target is fundamental at this stage. ML models work better when focused on a specific problem. You may find it difficult to make models perform well later when your targets are much too different from each other.



REMEMBER

A solution that works in European markets won't necessarily work the same in U.S. or Asian markets. Consumer culture and competitive scenarios are different. Seldom will a single algorithm grasp extreme variety, especially when your team is unaware of the differences and treats all the data the same.

Stating your problem and the elements compounding it helps your team extract the right data, process it correctly, and frame the best type of model for it. More importantly, understanding the problem helps your data scientists determine the correct error measures to track. Using this approach may sound a bit too esoteric because error measures are mathematical in nature, so that only experts in your team can deal with them, but if you don't explore the error measures, you won't know whether the model output is correct. It's essential to discuss which error measures you should use with the more expert among the data scientists in your team.

Keeping Up Trust and Interpretability

You must also realize that it's not just a matter of doing things, but of doing them in the right way. Your solution affects company prospects and determines how your customers see you, especially when something goes wrong. The news is crowded with examples about how AI can be biased, unfair, and sometimes discriminatory: from object detectors that don't work with all people to chatbots that suddenly spew hate speech.

When designing a solution, your team should keep in mind that error performance isn't the only metric a model should be evaluated against. There are stringent regulations, such as the General Data Protection Regulation (GDPR) in Europe. You must also consider society's unwritten rules about the trust expressed by consumers and public opinion. After all, trust is the only reason that people accept things they don't fully understand in a complex society. You don't use an elevator if you don't trust the engineers that built it and the maintainers that keep it working properly. The same goes for ML solutions and for any AI automation you may devise.



TIP

Model interpretability is the answer: Such methodologies explain why a model makes a certain decision in a particular case and provides you with ideas on why a model generally works in a certain way. Be sure that your DS platform offers you model interpretability because you need it.

Preparing the Data Properly

As part of the process, you gather data in a certain quantity and of a certain quality using the principle of garbage in, garbage out (GIGO in short) as a guide. Bad data can harm your AI implementation. Missing data, errors, and noise in features hide signals in data. Anomalous values called *outliers* can skew results. These are all problems that your data scientists will have to handle properly, which can consume up to 80 percent of their time. This is the *data preparation* part.

In the end, you get a *data pipeline*, a series of procedures that start from your information sources and output the data in a form that you can use for ML.

Getting Into a Virtuous Circle

After you obtain and prepare the data, processing it using ML algorithms will start. Your data scientists use part of the data to build models — the data for the training phase or *training data*, and part of the data for their validation, the *test data*. They can also use a *cross-validation* technique to use the same data for both training and testing.

Even if experience guides your team on the best learning approach, the truth is that you can't know beforehand whether ML will prove effective for your problem. In addition, you can't know which algorithm will be the most effective. It's a matter of repeated trial and error, as your team iteratively experiments using different solutions, records their results, and tests new approaches to solve your problem. You can start from the simplest models and, as time allows, build more complex and effective solutions to the problem.



REMEMBER

Here you get the “science” in data science. As in any science, when working with data and ML, you test various hypotheses and, based on the results, you decide whether to formulate and test even more hypotheses.



TIP

Repeated testing of new ideas to reach higher performance is possible because you have an established data pipeline that you can easily modify and rerun. This is a clear advantage of DS platforms.

Testing Everything for Production

When you're satisfied with the tests of the solution your team has devised, it's time to leave the lab and put everything into production, which means integrating the solution into the daily operations of your organization. At this point, things become difficult because it's not enough that your solution runs on the DS platform — it must run in the complex production environment, which has different performance demands. For instance, now the solution will see use by thousands or millions of users. This is the *scalability* issue.

Moreover, you have to control the way you provide an update and, if necessary, fix solution problems while in production. Handing the update to the IT department to code from scratch (hopefully with a high grade of fidelity) and asking it to maintain the solution throughout time (often without the necessary skill set) isn't the best solution. In many situations, you won't have the potential to assign someone permanently from your team to oversee the data product.



TIP

API-based services allow for easy deployment, easy maintenance, and updating by your team without requiring much time and resources. You wrap the solution produced by your DS platform in an API, which makes it easy for you to scale it from hundreds to millions of users using the flexibility that only a cloud solution can offer.

Adjusting to Concept Drift

Even when you have your data product deployed in production and working, your task isn't finished. Market conditions do change. You soon have to deal with a different competitive scenario, mutated consumer preferences, a redefined portfolio of customers, and more technological innovations entering the scene. Given all that, it's not assured that your ML algorithm can keep up with the changes because it has acquired its predictive power given different conditions.



REMEMBER

ML generalization capabilities imply that its algorithms can extract general rules from data. However, data content depends on the manner and from which context you collect it. If data mutates, your ML model may prove ineffective. In such a sense, ML solutions aren't like traditional software that sits for years on your systems and never turns bad.

Consider a simple example of a model trained to predict house prices. House prices change all the time, and if your model is older than six months, it will most likely provide you with a wrong answer. The same goes for any model built on advertising or online shopping data. These sometimes slow, sometimes abrupt modifications in the context that your data represents are *concept drift*. There are few remedies to avoid this pitfall other than continuously monitoring the models' performance, and to periodically schedule both new data collection and retraining.

Concept drift requires that you reiterate the same processes repeatedly, using fresh data, and make the necessary modifications to processing and ML algorithms. Ad hoc solutions without a clear structure and a precise workflow are not easy to replicate, and you soon realize that you're mostly investing time and energy on reinventing the wheel.



TIP

Using a DS platform from the beginning is the only way to proceed if you want to continue to have a competitive advantage generated by data and ML solutions over time.

- » Using AI marketing success stories to your benefit
- » Explaining the role of the right tools as a key enabler
- » Recapping why you should invest in AI

Chapter 5

Learning from Leading Marketing Teams

In this chapter, you read of successful stories from leading marketing teams, and discover many other real life marketing AI projects that silently succeed and exceed expectations. All these stories have the same thing in common — they use the right approach and tools, while avoiding pitfalls.

Listening To Successful Stories

You should now see that AI-based methods in marketing are more effective than the heuristics of the past and that AI-based methods are a necessity in an evolving digital world. First, because of the massive amounts of data that AI can handle at scale and that humans couldn't otherwise analyze. Second, because AI is the new frontier of competitive advantage: Not pursuing AI in marketing at this time could lead to loss of market shares and the creation of a deep gap with companies that have made AI an integral part of their internal organization and orientation toward the market.

You should listen to stories about successful AI marketing solution implementations to determine whether you can learn any lessons from them. It's possible to find plenty of examples from

companies such as Amazon and Netflix. These huge companies own large data resources and employ hundreds of data scientists. Each company invested in AI and received significant returns, but they likely exist in a different reality from your organization. Finding stories about smaller companies and their less advertised successes in marketing AI could provide more inspiration, but are also harder to find. The following list provides you with a taste of successful AI stories in marketing for smaller companies:

- » **Stitch Fix:** Demonstrates how marketing AI can make an otherwise unfeasible business plan quite possible and shows how machine learning (ML) can deliver value directly to customers.
- » **McDonald's:** Shows the switch to an information-centric approach to assist in business decisions and optimize both point-of-sale and customer experience. This solution relies on a combination of data, ML algorithms, and multidisciplinary teams.
- » **Mazda:** Uses ML to identify influencers on social networks who can promote a new car model in the United States.
- » **Coyote (with the help of Dataiku):** Uses a predictive behavioral system to retain present customers, increase the use of their devices, and help acquire new customers.

With this overview in mind, it's time to review each company in more detail. Stitch Fix is a company offering online subscription and personal styling service in the United States. Initially limited to women's clothes, it has expanded to men's clothing, plus sizes, maternity wear, and kids. Its offer is simple: They send clothes and accessories to their customers based on their understanding of a customer's preferences and style. The company bases the item choices mostly on ML, turning their business from potentially capital intensive and risky to inventory fast and capital efficient because of the efficacy of their predictive system. Here ML isn't only supporting their business model; it's also a key part of it because without precise recommendations driven by detailed customer understanding, a business like Stitch Fix wouldn't be profitable and scalable.



TIP

Stitch Fix is a profitable company because it brings value to its customers through AI algorithms.

McDonald's recently started using more of its available data, combining multiple previously siloed data sources, to understand and better control distribution, and to enhance their customers' experiences. After setting up multidisciplinary teams of data scientists and business experts, they started by providing more insightful reports on restaurants using advanced analytics derived from different existing sources. Then they processed large amounts of ordering pattern data, point-of-sale data, video data, and sensor data to optimize in-store traffic, customer interactions, and the drive-thru experience.



TIP

Digitalization of menu orderings has enabled this ML revolution at McDonald's.

Marketing AI can also enhance and automate classic activities such as selecting a target for a marketing action. The car manufacturer Mazda has leveraged ML to scan through multiple social networks to find influential groups of people who can act as ambassadors for the new car model launching in the United States. Finding the same target using a search conducted by humans would have taken a long time and a lot of resources, but AI automation can achieve better results in a shorter time and with much less effort.

Coyote, the French leader of real-time road information, depends on its customer base for the richness and promptness of information it provides. The customers themselves provide the traffic information that Coyote then shares with everyone. For this reason, Coyote not only keeps profitable customers, but also acquires new ones. Automatically analyzing behavioral data using its connected device, Coyote can model user behavior leveraging ML and successfully estimate the probability of a customer abandoning the service. It uses the same technique to target new customer prospects (increasing the performance of outbound call campaigns by over 11 percent for new customer recruitment).

Understanding the Secret Ingredients

When talking about successful examples in marketing AI, recommender systems always take the spotlight. After all, everyone scrutinizes them, both the customers that use their suggestions and the competition that monitors how effective they are. Of course, the competition wants to determine whether it's possible to replicate the recommender in their business model.

Unless promoted by vendors who want everyone to know about their contribution and products, many AI marketing projects aren't advertised as much as the most popular recommender systems. That's because the results of such projects don't receive the notice of anyone but insiders. For instance, an attribution model is only used internally by the marketing staff. Companies prefer keeping their most valuable projects reserved, since these systems power their competitive advantage.

Consequently, you won't hear much about successful AI marketing projects in conferences or read much about them in articles and whitepapers. From my direct experience, I can tell you that all successful AI marketing projects distinguish themselves through clear and focused objectives, their adoption of technologies that scale, and their efficient execution. In particular, the adoption of a good data science platform is what discriminates most of these projects from the rest, since a suitable data science platform provides the means for a data team to be more productive and innovative through:

- »» A faster way to clean and wrangle data (efficiency)
- »» Coders and noncoders alike contributing meaningfully (team and project coordination)
- »» Clear, reproducible workflows and a way for team leaders to monitor data projects (reproducibility)
- »» Automation to alleviate the inefficiencies of rebuilding and retraining models (automation)
- »» Efficient means to deploy data projects into production quickly (deployment)

Chapter 6

Ten Things You Need to Know About Marketing AI

This chapter discusses some not-so-obvious suggestions to succeed with your marketing AI project:

- » **Don't aim for AI; aim for machine learning:** AI is getting all the hype these days, but building AI is complex and requires putting together too many resources and competencies. Aim first for a machine learning solution, mapping data to results. It's easier and more effective for your marketing project.
- » **Start with low hanging fruit:** When deciding on goals for your first marketing project, decide based on the data available and on business impact. A quick success creates momentum for your data science team. It's better to develop a churn model that saves customer acquisition money quickly, than to take years to unveil a complex customer behavior model.
- » **Most of the hard work is obtaining and cleaning the data:** The truth about many projects is that mathematically perfect learning algorithms matter less than data carefully prepared by your team's hard work. A good data science platform may actually prove more valuable than some fancy machine learning package.
- » **Feature engineering is the key:** Acquiring the data, putting it together, and cleaning it from evident problems is just the start.

To succeed, you need to elaborate the data and create features that don't exist initially. To use a house price example, starting with data about house surface and market price isn't enough. You must also compute price per square foot to give your machine learning solution a boost. Leverage your marketing knowledge to create meaningful features for your projects.

- » **More data beats the clever algorithm:** When you have lots of data, some problems become solvable whether you use simple or complex algorithms. If you have resources, invest them on getting more data — not just of the kind you already have but also incorporating third-party data (such as social media data, weather data, and so on).
- » **Try many models, not just one:** Even when the first algorithm you try offers promising results, spend time trying everything. You may be surprised to find something that works even better.
- » **If you have enough data, deep learning opens a world of possibilities:** Even if it requires more effort and use of specialized teams, deep learning can solve problems in a better way than any other algorithm when you have the right amount of data. For example, Voyage Privé, a luxury online travel agency, used deep learning to analyze vacation photos and determine which one was most likely to be clicked by a particular site visitor, thus providing recommendations that are even more effective.
- » **It is generalization that counts; don't fool yourself:** When you start, you have no guarantee that a machine learning project will succeed. After evaluating the results, use the scientific way to check the results with a test set or cross-validation. If you can't solve a problem with more data, it's better to acknowledge the difficulty and to move to the next problem, which might be successful.
- » **In a competitive marketing scene, models are often short-lived:** Even if your models are successful in the present context, they may soon become obsolete if you stop gathering fresh data and you don't update their parameters. The best practice is to establish a reusable and efficient data pipeline, using a data platform, to check your models' performances often and to be ready to update them periodically.
- » **Don't believe in vendors promising the silver bullet:** At a certain point, some vendor will promise a magic algorithm capable of predicting everything. Don't believe that, please. Nothing can replace hard work on data, and there is no hidden magical algorithm in a world of open source solutions.

Your Path to Marketing AI

Marketing use cases are often the lowest-hanging fruit and can unlock the path to Enterprise AI for the entire company. That's why today's top marketing teams are leading by example by diving in to machine learning, proving its worth, and thus leading the AI revolution.

More than 3,000 marketers at some of the world's most progressive companies are using Dataiku to lead the organization in its AI progress.



SEPHORA

FOR EVERY USE CASE



**Churn
Prediction**



**ML-Based
Attribution**



**Pricing
Optimization**



Forecasting



**Hyper-Targeted
Advertising**

FOR EVERY INDUSTRY



**Retail
E-Commerce
& CPG**



Pharmaceuticals



**Banking &
Insurance**



**Media &
Entertainment**

AND FOR EVERYONE:



**Data
Scientist**



Analyst



**Team
Manager**



**Marketing
Professional**

[learn more](#)



Bring AI-focus to traditional marketing analytics

Data science, machine learning, and artificial intelligence have changed the nature of marketing. Companies and marketers need to understand these tools so they won't be left in the dust by well-executed AI marketing projects from competitors. This book explains why this revolution is happening today and gives you information on key AI and machine-learning topics.

Inside...

- Shift from heuristics to AI-based solutions
- Get a clear idea of key AI concepts
- Understand the use cases for AI in marketing
- Get started with a marketing project using AI
- Avoid common pitfalls in your AI project
- Read success stories from marketing teams



data
iku

Luca Massaron is a data scientist who specializes in organizing and interpreting big data and transforming it into smart data. Together with John Paul Mueller, he coauthored *Machine Learning For Dummies*, *Artificial Intelligence For Dummies*, and *Deep Learning For Dummies*. He is a Google Developer Expert (GDE) in machine learning.

Go to **Dummies.com**®
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-119-59706-3
Not For Resale

for
dummies®
A Wiley Brand



Also available
as an e-book



WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.