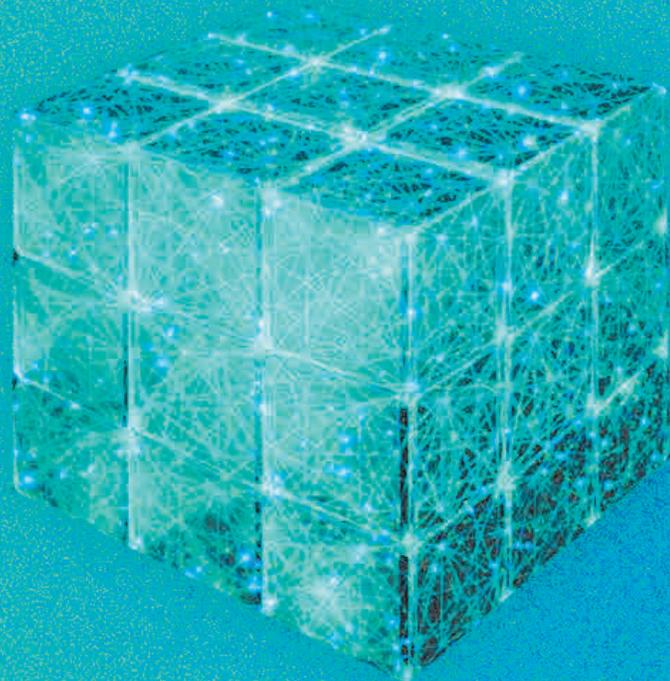


SETH EARLEY

THE  
AI-POWERED  
ENTERPRISE



HARNESS THE POWER OF ONTOLOGIES  
TO MAKE YOUR BUSINESS SMARTER, FASTER  
AND MORE PROFITABLE

This is a chapter excerpt from *The AI Powered Enterprise* - "Making Ecommerce Smarter".

I would love to get your feedback on the content. Also, for a limited time, we are offering complementary copies of the book to people who will provide an Amazon review and inform their social and professional network about *The AI Powered Enterprise*.

For a complementary physical copy of the book, contact [Carolyn.Southwick@earley.com](mailto:Carolyn.Southwick@earley.com) and let her know you will post a review and how you can help build awareness.

You can also reach me at 781-820-8080 or [seth@earley.com](mailto:seth@earley.com) to share your feedback or ask me questions. Happy to chat about AI and IA!

Thanks for your support!

Seth Earley, Author of *The AI Powered Enterprise*

## Praise for **THE AI-POWERED ENTERPRISE**

“AI promises to provide the next ‘turn of the crank’ in business automation. However, purely statistical machine learning alone won’t achieve this on its own. This book provides prescriptive guidance in the context of real business case studies to drive success instead of disappointment. It’s a great resource to separate the hype from the reality and a practical guide to achieve real business outcomes using AI technology.”

—**Peter N. Johnson**, MetLife Fellow, SVP, MetLife

“AI and its various meanings bring a whole new dimension to how enterprises operate. In his book *The AI-Powered Enterprise*, Seth has honed in on the methodology behind how AI impacts content—the one constant tangible for every enterprise. A must-read for any enterprise interested in what their content says about them—from a data miner to a taxonomist to a casual blogger. The clarity of the various aspects of content design is super impressive.”

—**Eeshita Grover**, Director of Marketing, Cisco

“Artificial intelligence holds the power to transform your business and your career, but there will be plenty of challenges along the way. Earley demystifies the topic and provides a practical road map for applying smarter processes and technologies across the enterprise. Now is the time to explore AI, and this book is a great place to start.”

—**Paul Roetzer**, founder and CEO, Marketing Artificial Intelligence Institute and author of *The Marketing Performance Blueprint*

“If you wonder where to get started in your company’s AI journey, Seth lays out a playbook for a flexible, scalable, and tool-agnostic approach to AI implementation. The journey is difficult, but the rewards will enable businesses to compete and thrive in our complex and fast-changing technology world.”

—**Ryan Miller**, Chief Digital Officer, First Investors Financial Services

“With every technology that comes knocking on our doorstep, the question we have to ask ourselves is, ‘how will I use this to create value for my customers?’ Businesses have weathered the turbulence of digital transformation and now AI is knocking on our doorstep. For executives, it can feel confusing, intimidating, and at times, downright dangerous. If you’re in that mix, then you must keep Seth Earley’s book close at hand. It delivers a practical approach to explaining AI as well as how to apply it to any organizational environment.”

—**Carla Johnson**, international keynote speaker, bestselling author, and CMO

“For any leader considering ways to improve their business with advanced data analytics and artificial intelligence, this book is a must-read. Seth Earley has documented a recipe for your success.”

—**Mark Loboda**, Senior Vice President of Science and Technology, Hemlock Semiconductor

“Read this book to learn how leaders and companies are using AI with structured data to transform business. Insight from real-world examples, combined with a proven methodology, will arm the reader with the knowledge and confidence necessary to drive AI in any organization.”

—**Barry Coflan**, SVP & Chief Technology Officer, Schneider Electric, Digital Energy

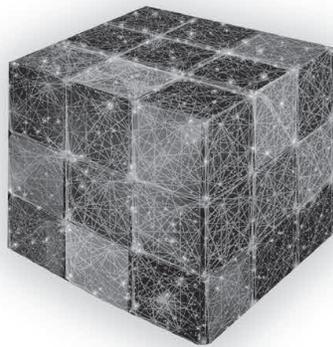
“Success comes to those who plan ahead. And while AI has been in existence for a while, it’s still not working at the level it should be. Why not? Seth’s book provides the answer to that question and a clear path for everyone who want to hang ten the AI wave. Don’t crash & burn, read this book!”

—**Jeffrey Hayzlett**, Primetime TV and podcast host, speaker, author and part-time cowboy

# THE AI-POWERED ENTERPRISE



# THE AI-POWERED ENTERPRISE



Harness the Power of Ontologies to  
Make Your Business Smarter, Faster,  
and More Profitable

SETH EARLEY

Copyright © 2020 by Seth Earley  
Foreword copyright © 2020 by Thomas H. Davenport

20 21 22 23 24 5 4 3 2 1

All rights reserved. No part of this book may be reproduced, stored in a retrieval system or transmitted, in any form or by any means, without the prior written consent of the publisher or a license from The Canadian Copyright Licensing Agency (Access Copyright). For a copyright license, visit [www.accesscopyright.ca](http://www.accesscopyright.ca) or call toll free to 1-800-893-5777.

Cataloguing data available from Library and Archives Canada

ISBN 978-1-928055-50-1 (hardcover)

ISBN 978-1-928055-52-5 (EPUB)

ISBN 978-1-928055-51-8 (PDF)

Cover and interior design: Greg Tabor

Cover image: iStock/antoniokhr



Published by LifeTree Media Ltd.

[LifeTreeMedia.com](http://LifeTreeMedia.com)

Distributed in the U.S. by Publishers Group West and in  
Canada by Publishers Group Canada

Printed and bound in Canada

# CONTENTS

Foreword / ix

Chapter 1: The Promise and the Challenge of AI / 1

Chapter 2: Building the Ontology / 27

Chapter 3: Customer Experience: The Front Line of the  
Battle / 50

Chapter 4: Marketing: At the Center of the Data Storm / 76

Chapter 5: Making Ecommerce Smarter / 87

Chapter 6: Sharpening the Sales Process / 121

Chapter 7: Customer Service: Delivering Higher Quality at a  
Lower Cost / 144

Chapter 8: Accelerating Employee Productivity / 163

Chapter 9: Physical Meets Digital: Manufacturing, Supply  
Chain, and Logistics / 189

Chapter 10: AI-Powered Strategy and Governance / 215

Chapter 11: Leading into the Future / 236

*Glossary / 273*

*Acknowledgments / 287*

*Notes / 290*

*Index / 296*

*About the Author / 305*

*To my wife Lisa, the love of my life,  
for supporting me throughout this process  
and for her support over the years  
of building the business.*



## FOREWORD

**W**hen I think of Seth Earley, I think of four things: knowledge exchange “salons” around the firepit at his home, a crazy-fast and small racing-power catamaran on which I risked my safety with him, knowledge management, and conversational artificial intelligence. I won’t elaborate further on the first two, but the last two are highly relevant to this book. Like me, Earley spent a number of years consulting and writing about knowledge management. That topic isn’t terribly popular now, but it has positively informed his approach to AI and this book.

The book is heavily focused on the importance of ontologies in conversational AI systems, and ontologies were also very important in knowledge management. Of course, the ontology Earley is addressing is not the metaphysical study of the nature of being. The other meaning of ontology is, to use Google’s definition, “a set of concepts and categories in a subject area or domain that shows their properties and the relations between them.” Earley is convinced, as am I, that you can’t create an effective conversational AI system—a chatbot, an intelligent agent, or a virtual assistant—without an ontology. And if you want a practical introduction to ontologies and their application in conversational AI, you have come to the right place.

I have always liked Seth’s favorite line of, “There is no AI without IA,” where “IA” is information architecture—another term for ontology. It’s both catchy and mostly true. Of course, as an academic (though a pretty

practical one), I am inclined to look for exceptions to general rules, and I think there are some exceptions to that sensible rule.

You will note in this book that the primary examples of AI that Earley uses are the conversational approaches that organizations are widely adopting for interactions with customers and employees. His appealing rule about AI and IA is true for the great majority of conversational AI applications. Perhaps there are exceptions to it if you include, for example, statistical translation programs like Google Translate in the category of conversational AI. I'm not sure they belong in the category, but they do facilitate conversation with speakers of other languages.

The other exception domain is also debatable. I have argued with Seth that machine learning—statistical approaches to prediction and classification that many observers include as AI—doesn't necessarily demand an ontology. It requires data (often great heaping gobs of it) on multiple variables (also called features) in rows and columns, but we don't need to know the terminological relationships between them.

But Earley's counterargument is that if you are managing enterprise data, it is hugely beneficial to have an ontology that describes what data entities you have in your organization and how they relate. If you have a customer ID data element, for example, it's very useful to know that it includes current customers, former customers, and prospects, and that customer orders are aggregated by that ID. Earley has a particular passion for customer and product data, and I agree that they are perhaps the most important types for most organizations.

I am not a fan of massive projects to create enterprise-level ontologies (often known as “master data management” or “enterprise information architecture” projects), but there is no doubt that your organization should know what data elements you have access to, and how they relate to other data. And I believe that ontology efforts are of more value if they are relatively narrow in scope. So I partially concede the argument to Earley even for machine learning. And in the book he describes a “bottom-up, data- and content-centric” approach to creating ontologies, which I generally believe is more effective.

In any case, there is far more to this book than urgings about ontology. It also addresses such topics as the critical role of data in digital

transformation, techniques like tagging data with metadata, organizational challenges for effective data, and the role of data in the customer experience. My favorite story in that regard involves granite countertop cleaners, but I won't spoil your suspense by relating it.

A great strength of the book is that it is replete with such examples from Earley's experience and the work of his consulting firm. He has worked with companies across many different industries—many of them very large and successful firms—that had problems with their data in the context of AI. His firm has built successful and sophisticated intelligent agent or chatbot systems for both customers and employees.

I have written two books on AI and read many more, and I do not know of any books that have such useful and detailed advice on the relationship between data and successful conversational AI systems. If that's what you need, read on.

*Thomas H. Davenport*

*President's Distinguished Professor at Babson College, Research Fellow at MIT Initiative on the Digital Economy, and author of Only Humans Need Apply and The AI Advantage*



Marketing is about more than messaging as it expands to include collaboration and engagement. Customers want to have their problems solved, not to be marketed to. This may mean exposing the knowledge and expertise of engineers in a B2B organization, or it may mean enabling customers to have the insights they need to choose your product in a B2C context. This means that marketers have to be knowledge enablers; they are no longer limited to the role of messaging creator.

Marketing is now about scaling the machinery of communication, collaboration, and content processes, much as Molly Soat did at AMA. It is about enrolling deeper levels of the organization in the process. This means marketers must be involved in various aspects of governance and change management to get meaningful content created, managed, organized effectively, and presented to target customers in a consumable way. Marketers also need to be intimately involved in IT processes. They need to work closely with the CIO (chief information officer), and, if the role exists, the CDO (chief data officer). Marketing is increasingly data-driven, and data quality will be essential to marketing success. Digital agility means data agility. Therefore, marketers need to show the organization what can be done with high-quality data as well as demonstrate the impact of poor quality, difficult to access, or missing data. Marketing leaders can champion data quality to show the impact on the bottom line through metrics linked to customer acquisition and revenue growth.

I recently heard about an organization choosing a digital asset management system *without* the involvement of marketing. Marketing is one of the chief creators and consumers of non-text assets, and it should therefore be driving requirements for these types of systems, as well as systems including web/marketing content management, marketing automation, and customer relationship management (CRM).

Ultimately, digital marketers need to become digital architects. The marketing function leverages data assets from many parts of the enterprise—from customer purchase histories to call center feedback, survey responses, social media data, clickstream behaviors, campaign responses, external data feeds, mobile usage data, and search metrics. Deriving value from these sources means translating this “digital body language” into meaningful content, campaigns, and offers. Increasingly, this means

translating data models from various systems into attributes managed in the ontology. Those attributes become inputs into personalization engines, web content management tools, collateral creation processes, campaign management systems, and various outbound demand-generation activities.

All of this leads to one conclusion: the ontology is very much the responsibility of the senior leaders in digital marketing. The digital marketing landscape is constantly changing, with thousands of vendors clamoring for attention to their “AI” technology. There is no shortage of white papers, books, and conferences on marketing technology, customer experience, or the analytics of engagement. Executives are getting advice from all quarters about what they need to do to have a successful digital marketing strategy.

What is conspicuously missing from these to-do lists is any reference to the foundational role of ontologies. Many systems are deployed in isolation or with a nod to integration using web services; however, very few marketing leaders are in a position to develop the foundational data infrastructure that is needed for success. If the enterprise is to have any hope of a positive outcome from all of the investments being made in advanced marketing technologies that are meant to smooth the customer journey, marketing leaders have to streamline marketing operations and supporting processes across all of their tools. They must pursue this effort in a holistic way that includes a framework so all of these systems can communicate.

The next chapter will provide a deeper dive into ecommerce and how maturity in product information and content processes will arm the organization for the ongoing battle for ecommerce wallet share.

## TAKEAWAYS FROM CHAPTER 4

In this chapter, I've described how digital marketers play a central role in integrating customer-facing technologies and must connect them to a central ontology. These are the main points in this chapter:

- Marketing's job starts with awareness by delivering the appropriate piece of content at any moment to move the user to the next step in their journey.
- Agile digital marketing requires orchestrating messages across diverse systems, not all of which are controlled by marketers.
- Digital marketers must become knowledge enablers, champions of data quality, architects of digital systems, and keepers of the ontology that powers it all.



## CHAPTER 5

# MAKING ECOMMERCE SMARTER

**E**commerce used to be about finding what you wanted to buy. But there are so many possible things to buy—and so much information available about the people who want to buy them—that ecommerce has to be smarter than that now.

What does smarter ecommerce look like? Let's imagine, for a moment, that my wife is buying me a pair of hiking boots . . . and that she's interacting with a virtual sales assistant on a website. Virtual sales assistants can guide consumers in locating and selecting products and options, depending on a range of criteria. One method is faceted search, which allows the customer to narrow the selection according to criteria such as brand, style, color, price, and size. Here's how that conversation might go:

Virtual sales assistant: "What type of hiker is your husband—serious multi-day backpacking, or day trips?"

Her: "Day trips."

VSA: "What length boot does he prefer—ankle or calf?"

Her: "Ankle."

VSA: “Any special features?”

Her: “Very wide feet.” [I have been told I could probably water ski without skis.]

VSA: “That limits some other choices. What type of weather?”

Her: “All seasons, I think.”

This conversation could continue with more questions: Do you need something waterproof or insulated? Lightweight or heavy? Is there a price range you need to stay within? With a dialogue like this, the virtual assistant narrows the choices down to the exact boots she should buy for me.

Behind the scenes, powering this virtual sales assistant, is structured data about the features that are important to customers. This chapter is about how to prepare that data and tap artificial intelligence to enable a more interactive, customized, conversational approach to ecommerce.

## ECOMMERCE IS WHERE AI HAS ITS BIGGEST IMPACT

As I described in chapter 3, customer experience is about the end-to-end process of acquiring, serving, and supporting the customer. It spans nearly every department and process in the enterprise.

Ecommerce, websites, and apps are a narrow slice of that experience, but that slice is crucially important. It’s the most powerful place to improve the customer experience. It’s connected to all the company’s data systems, and even in companies from which people don’t typically buy online, online is often where they make their buying *decisions*.

Even though ecommerce is just one aspect of the customer journey, it generates revenue directly. As a result, you can use it to justify investments that improve all aspects of the customer experience. Change an ecommerce site—whether through a redesign, new technology, changes to taxonomy, improvements in supporting content, or whatever—and you can directly measure the results. This means that ecommerce can create the foundation for metrics-driven governance, the decision-making playbook that is the cornerstone of a data-driven organization. (I’ll discuss governance in more detail in chapter 10.)

In this chapter, I will share an in-depth exploration of the data,

architecture, governance, and execution elements of successful ecommerce and interactive programs. I will refer to the web presence as “ecommerce” whether the site is transactional or not. Regardless of whether customers actually purchase on a site, design and execution are crucial when educating the customer about product features, specifications, availability, and where to purchase (if not online).

Here’s how I’ll tell that story. First, I’ll show how search is giving way to smart, conversational, predictive forms of interaction as a way to help customers, based on solid product and customer data. Then I’ll show you how such data can lead to highly personalized and customized sites that connect each customer as quickly as possible to their best possible offer, powered by the ontology. And finally, I’ll describe how you can assess your tools to determine if your business is ready to deliver on the promise of smarter ecommerce.

## HOW AI TURNS SEARCH AND SITE NAVIGATION INTO A CONVERSATION

Questions are ambiguous. To understand them, you need to know more about who is asking and the context.

For example, I once attempted a guided winter hike up Mount Washington—the highest mountain on the east coast of the United States—on the coldest day of the year. When the guide asked me about my experience, I inferred from the context that he was asking about my experience with winter mountain climbing, not my experience as an AI consultant. Similarly, in a business scenario, if I don’t know anything about the person asking a question (perhaps at a networking event rather than during a sales call), I might ask a clarifying question in order to answer more precisely.

Search is now a conversation as well. Search terms are often short, ambiguous, and an approximation of the searcher’s real information need. A generic set of responses is no longer the best choice. When you start to type a search into Google, the query box will begin to contextualize the query with suggestions and to customize the results based on your own personal search history. Why? *Because search is a conversation.*

On an ecommerce site, the same sort of thing happens. For example, if you visit MSC Industrial Supply's website, [mscdirect.com](http://mscdirect.com), and search "tools," the site will try to anticipate the category of tool by presenting choices that clarify the question, What kind of tools are you looking for? If you select "power drills" from the set of resulting options (also known as facets), the site will then narrow down, attempting to determine "What kind of power drill?" by presenting a selection of air drills, cordless drills, coring drills, electric drills, and so on. Select "electric drills," and the site then narrows things down further by determining more details: brand, size, speed, handle, and so on. These clarifying choices appear as navigational elements. We are, you might say, drilling down into drills to get to the specific item that you need.

On-site search is a critical function on an ecommerce site, and one in which AI techniques can balance recall (the completeness of results) with precision (the most appropriate results). Recall and precision are typically at odds with one another, but AI that has the correct information architecture can increase both recall *and* precision. For example, a smarter search engine recognizes that when I search for "sushi," the results should match restaurants close to my location.

A smart site needs context to improve results. Consider a customer searching an industrial site for "mold stripping." They may be a plastics manufacturer who wants to clean an injection mold, an industrial reclamation contractor seeking to remove mold and mildew from damp surfaces, or a remodeler resurfacing the wooden molding on a ceiling. Depending on the customer's context, the correct search results might be release lubricants, cleaning chemicals, or abrasives. The search term itself is only one signal; add the industry or application and you have supplied another signal that the search engine can interpret to more precisely return the appropriate result.

Where does context come from? Each dimension that informs a search is part of the ontology: classifications of products and customers, the desired application, and the industry, for example. The search engine uses relationships contained in the ontology to contextualize industry and return the correct products.

Search has long depended on AI and machine learning techniques to

contextualize vague or ambiguous queries. AI enables functionality like analyzing text, clustering and grouping content, creating search indexes, extracting entities from text (such as name, company name, dates, or part numbers), interpreting intent, correcting spelling, and presenting related items. AI can inform search results by weighing other signals about the customer's context, including whether or not the customer is new, technically sophisticated, or commercial, and what tasks, problems, objectives, and interests they may have.

Each of these concepts is modeled in the ontology, which contains the relationships that are important to the customer, industry, and engagement strategy. In the case of the “tools” search term, for example, Amazon, Lowe's, Grainger, Gamut, and MSC all return different results and have different conversations with the user. An intelligent assistant can convert these queries into a conversational interaction:

User: “I am looking for tools.”

Bot: “What kind of tools do you need? Power tools or hand tools?”

AI can take the various phrase variants (“I need tools,” “I am looking for a tool,” and so on) and interpret the customer's intent based on historical data (search logs or chat logs, for example) to return the correct tool categories and continue refinement from there.

### **Dynamic Navigation**

Search and site navigation are two sides of the same coin; in a smart site, both processes improve incrementally based on context. In the ecommerce facet refinement scenario that I just described, each of the clarifications looked like navigational elements: links or checkboxes. The customer may have felt as if they were navigating, but behind the scenes, the system was executing a search. The navigation is dynamic, in that the order and detail of choices surfaces the most appropriate attributes and characteristics based on other signals from the user.

How should the questions the bot asks, or the options it presents, be ordered? What is the correct level of detail to show the user? A site can order navigational links based on rules (bestselling product categories first, perhaps, and then highest margin categories) or based on customer knowledge. If you know that this customer has always purchased sale

items, you might display the best deals or clearance categories first. A navigation rule could vary what is shown based on what items the customer has previously clicked on or purchased. For example, if a customer browsed running shoes and then purchased outdoor wear, a subsequent search for shorts could present activewear at the top of the product listing.

Sites can go too far with this level of customization. Ever-changing navigation breaks the customer's mental model; people want a sense of the structure and organizing principles of the site, so when things are moved around or the same category is used in more than one place—a so-called polyhierarchical structure—the customer may have trouble grasping the site's organizational structure. (As an analogy, when you go to the hotel's breakfast buffet, you might find the spoons next to the cereal, next to the coffee, or with the other silverware—but are they where you expect them to be?) Polyhierarchy can reduce findability and increase the customer's cognitive load, as there is no longer a simple answer to the question, "Where would I find this thing I'm looking for?"

Whatever rules your site uses—according to merchandising priorities, the needs for a given customer, what's overstocked, or a combination of factors—those rules depend on an intentional design for elements captured and referenced in the ontology. An ontology containing customer interest categories can serve them up to the search engine when that customer comes back to the site, for example. AI search tools and mechanisms are endlessly configurable, but to be effective they require fine-tuning, an understanding of the customer, and an information architecture that contains the appropriate product features, functions, and details that are important to target customers. You should build the appropriate ontology to drive search over time, rather than purchasing a generic or off-the-shelf architecture or engagement approach. Search personalization isn't one size fits all.

### **Predictive Offers and Shopping Basket Analysis**

Let's look at two of the most common mechanisms for presenting products and content that best serve a particular customer.

*Predictive offers* attempt to induce behaviors by anticipating a buying signal from a prospect and nudging them into a purchase. By testing the

performance of similar offers on similar customers, the algorithm can fine-tune offer parameters and design elements to increase likelihood to buy.

*Shopping basket analysis* makes a suggestion based on what others have purchased along with the products that the customer has selected. This can be refined by comparing shopping baskets of customers who are most similar to the customer currently shopping, who share the same industry, who have similar problems, or who are in similar job roles.

The ontology is the reference point for defining these characteristics consistently across systems. These types of offers can be surfaced in the context of a user's search or presented as navigational choices. In either case, they work in similar ways—interpreting user intent and combining that intent with other signals to provide a targeted response. That response could be a recommended product, a more functional (and expensive) option, a related accessory, an offer of some sort, or a solution. Over time, algorithms can learn from customer behavior and present content and offers that are progressively more effective.

## DATA POWERS ECOMMERCE

Since the ecommerce customer experience is made up completely of data, that data needs to be correctly designed, organized, and managed. While this sounds obvious, many organizations have immature product information processes. When they add new products (a process called “onboarding”), they don't manage product information in an adaptable, sustainable way.

An effective ecommerce experience begins with good product design principles and clean, well-structured product information, and then attempts to reduce the customer's cognitive load in navigating the site. What appears to be a simple customer experience on a site is, behind the scenes, quite complex. In fact, the role of the designer is to hide the complexity from the customers by thinking through the steps of their processes and the details of their mental models and presenting information in a manner meant to be intuitive. The information then makes sense to the customer. It is easy for the customer to find what they need, because

the designer anticipated their need and provided a limited number of choices based on that need and the customer's journey stage.

The designer begins this process manually and uses AI technologies to help refine, tune, and scale individual experiences. This requires a deep understanding of customers and how they think about the world. AI-powered ecommerce allows for differences in how users perceive the world, and should be designed to adapt to their needs based on real-time signals and behaviors. However, the process must begin with a foundational design that the AI can adapt and improve.

The fuel for a customized ecommerce experience comes from two kinds of data: product data and customer data. Let's start by looking at product data and taxonomies.

### **Product Taxonomies Are the Foundation of the Ecommerce Experience**

Managing a large product selection of thousands or even millions of products begins with a primary hierarchy called the product taxonomy. This is also called a navigational hierarchy or display taxonomy, and it is a critical foundation of the user experience. At first glance, one might think, "How hard is that?"

In fact, a number of years ago, when I was meeting with a new primary care physician, he asked me what kind of work I did. When I explained what taxonomies were—a way of organizing and navigating products, content, and knowledge—he responded with an incredulous, "You make a living at that?" I replied that I (along with 30 employees) indeed made a living at that. He could not comprehend why that would take any effort or how that could be a business. Similarly, many senior executives cannot understand why taxonomies would be costly or complex to develop or maintain.

Here's one way to think about the value of a taxonomy: The way products are organized within it is analogous to the way a physical store is designed. The taxonomy undergirds the virtual store and becomes part of the character, personality, and differentiation of your site.

Physical stores selling the same merchandise have different ways of presenting it. Products are arranged in different ways and merchandised to

appeal to certain types of customers. The taxonomy is like virtual store shelving (and aisles and signage). Different users not only shop for different items, but they also walk through the store to navigate and find their way differently.

Retailers spend a lot of time and money designing store layouts and build “planograms” to show where products should optimally be displayed and merchandised. By powering the ecommerce site, the product taxonomy functions as a virtual planogram that can, unlike a physical store layout, adapt to the needs of a customer as they move through the site.

For large ecommerce websites, product arrangements need to be tested across different customer types or personas (as I explain later in this chapter). Machine learning can then adjust and fine-tune assortments, offerings, and arrangements, but these should begin with a well-designed and tested hierarchy: the product taxonomy. Your product taxonomy is not necessarily going to be the same as that of your competition, even if you carry the exact same merchandise. The foundation has to be tuned specifically to the needs of your users and adapted to their circumstances, goals, and objectives.

Another challenge is in the quality and design of product data, which is directly related to customer needs. Product data should reflect the product features that are important to different types of customers. If I like socks with thick, fuzzy soles, but that attribute was not modeled in the product architecture, I won’t be able to find those types of socks as easily. Including a feature in the product data also enables analysis of purchases to inform promotions. I can send fuzzy sock offers to people who have purchased them or who are similar to people who have purchased them. A shallow understanding of the customer will impact the way that product data is architected and populated, resulting in missed opportunities.<sup>1</sup>

### **Product Relationships Must Be Included in the PIM System**

PIM (product information management) systems hold information about products, including product relationships. Product relationships establish connections between products, helping users navigate and discover products and solutions. Product relationships may be simple accessory relationships or may reflect complex component assemblies and

configuration rules. These relationships enable recommendations based on what products work best together or are used together as a solution. Recommendations about related products lead to higher revenue and more frequent returns to the site.<sup>2</sup>

Solution bundles can become quite complex; merchandisers, solution engineers, or subject matter experts may need to define them manually. But they can also be created using AI tools to mine information from maintenance manuals, product configuration content, and engineering specifications. Various product relationships can be captured and managed for presentation to users in the right context, including replacement parts, related parts and categories, kits and sets, competitor cross-references, distributor cross-references, compatible products, required components, optional accessories, and obsolete products. Your site can present these relationships to help the user choose what combination of products to buy, smoothing their journey.

Product relationships often already exist, but they need to be captured and normalized in a PIM. Product relationships can be discovered in static text, spreadsheets, folder systems, or existing PIM systems. Product managers and subject matter experts can help curate such product relationships. Understanding these relationships will help customer experience designers include the correct relationship for a particular use case. Best practices for managing product relationships include centralizing product relationships data, capturing and documenting the definitions of relationship types for consistency, creating processes for ongoing maintenance and governance, integrating relationship identification into the item onboarding process, and reporting regularly on product relationships by type (including usage metrics).

### **Standardize for Efficiency; Differentiate for Competitive Advantage**

There's a fundamental question when it comes to product taxonomies: should we match industry standards or develop our own customized version?

A key principle here is that standardization leads to consistency and efficiency, whereas differentiation leads to competitive advantage. As I have described, the navigational taxonomy informs the front-end customer

experience and guides how customers find the desired product on the site. Within your company, the taxonomy also influences how your merchandisers plan inventory and promotional strategies and how the finance organization predicts demand and tracks revenue and profitability.

In some cases, a manufacturer will want to use the same terminology, organizational structure, and navigation as their distributors—which may also mean that they are using the same conventions as their competitors. The upside is that this makes it easier for an organization to exchange information with trading partners and consistently track information as it flows through the supply chain.

Of course, if every competing site is based on a similar taxonomy, then a business needs to differentiate its products in another way. Ecommerce sites can still differentiate their offerings through such factors as price (larger sellers can offer bigger discounts), quality (premium products), or convenience (faster delivery, easier transactions, more intuitive site design, better suggestions, or more customized interactions).

### **Customer Scenarios, Personas, and Models**

The taxonomy of product features is, of course, only half of the equation when it comes to an organization delivering the best possible site experience. Data about customers is the other half.

How do we gain a deep understanding of the needs of our customers and tailor their experiences appropriately? We do so by modeling our users through the use of *personas*. Personas are iconic representations of the different types of customers that we serve. Customer experience teams typically develop these representations with input from other departments that have customer knowledge and systems of record. Some designers define personas at a high level without personal details (some refer to this higher-level classification as an “audience”). Such broad personas have names like “first-time buyer” or “repeat customer.” Other designers get into significant details about the wants and tastes of specific classes of customers.

Personas, while fictitious, can be developed with a lot of personal details that are meant to help the designer keep customer needs and preferences in mind. However, it is easy to get carried away and spend too much time on detailed persona development without spending enough time on

what the personas are *trying to achieve* through customer stories. A solid understanding of what personas care about and value comes from an often undervalued activity: primary research interviews with actual customers. When executives and merchandisers imagine that they know their customers and how they interact with the organization, their teams may develop personas without costly in-depth customer interviews and research. They may make unquestioned assumptions about their customer, resulting in hypothetical personas that can turn out to be inaccurate.

When personas are developed through actual research, however, the results can be eye-opening. Consider an example persona: Liz, a procurement manager with a graduate degree and two high-school-aged children. (Some teams choose names as hints about the customer—for example, if Liz is price-sensitive, she might be called “Low Budget Liz.”) Another persona might be Tom, a manufacturing line manager in his mid-fifties who likes to play golf and travel with his wife now that the kids are out of the house. Perhaps he is more technical and interested in quality and reliability—which is why he’s “Techie Tom.”

Modeling a persona puts the designer in the shoes of different types of customers, so they can understand those customers’ habits, goals, values, and ways of thinking. The persona informs design for their specific needs. Personas with nuance—a name, hobbies, a family, a career, particular values, and a personality—help designers better understand the types of target customers the organization is focused on. Then the people who are designing the site experience can ask themselves, “Would Tom think like this? How would Liz react to this offer?”

Developing a persona involves the creation of what are variously called customer stories, user stories, scenarios, or use cases. A customer story gets into the details of common challenges and problems across various types of users. The stories can get more granular as the needs and experiences of different customer types diverge, including increasing detail about the problems they are trying to solve and the tasks they need to complete on the site. As user experience expert Kim Godwin said, “Personas without scenarios are like characters with no plot.”<sup>3</sup> Personas and scenarios can also ensure that testing covers the range of customer types that are likely to use the site.

## Testing Across the Customer Taxonomy

A large, diversified business serving a range of markets and types of buyers could have numerous audiences and personas defined by characteristics like background (business or technical), role (buyer or engineer), title (CIO or CDO), size of firm (over or under \$1 billion in revenue), or buying authority (over or under \$1,000). Personas can help more precisely define audiences and inform design and testing if there are significant differences among audience members that impact navigation, terminology, goals, or other aspects of the customer's mental model.

Whatever you call them—audiences, personas, or customer types (and we'll use the terms interchangeably from here on)—designers use these classifications to make taxonomy and customer experience decisions. During taxonomy and customer experience design, personas, along with their scenarios, determine the information needed at a particular stage of their experience. Designers recruit test participants to ensure that the design meets the needs of different types of customers. Personas may not cover every type of customer (there are always outliers), but they should be representative of a cross-section of the majority of your customers.

Use multiple personas to test a core design on the different mental models of your customers. Different types of customers are likely to use different terminology, which should be captured and represented in the ontology. The ontology can act as a Rosetta Stone for your customers and is especially valuable when a website is serving both a technical/professional audience and a lay audience. For a health care site that serves both physicians and patients, for example, the patient search term for “cardiac care” might be “heart health.” The ontology must be responsive to both types of searches to meet the needs of both customer groups.

## Terminology, Context, and Expectations

While this may seem to be splitting hairs, it is important for an organization to test both the terms and concepts that are represented (the taxonomy) and how they are presented (the customer experience). They are intertwined, and therefore both must be tested.

When serving many audiences, the same terminology can have multiple meanings and contexts (remember “mold stripping”?). Consider a search

on a financial services site for “tax planning.” The right answer depends on the customer’s job: a financial planner may be looking for tax-free bonds, while an accountant is seeking detailed tax code rules governing financial instruments. Perhaps the site has a broad audience; it serves young people, older folks, and people of various demographic backgrounds and interests. What would be the point of describing all of these types of customers—is the experience really going to be that different for each? Instead, test performance across personas to identify differences, if any, in how people navigate. It may make sense to allow them to self-select (commercial versus individual) and then provide different terms, paths, and structures for each.

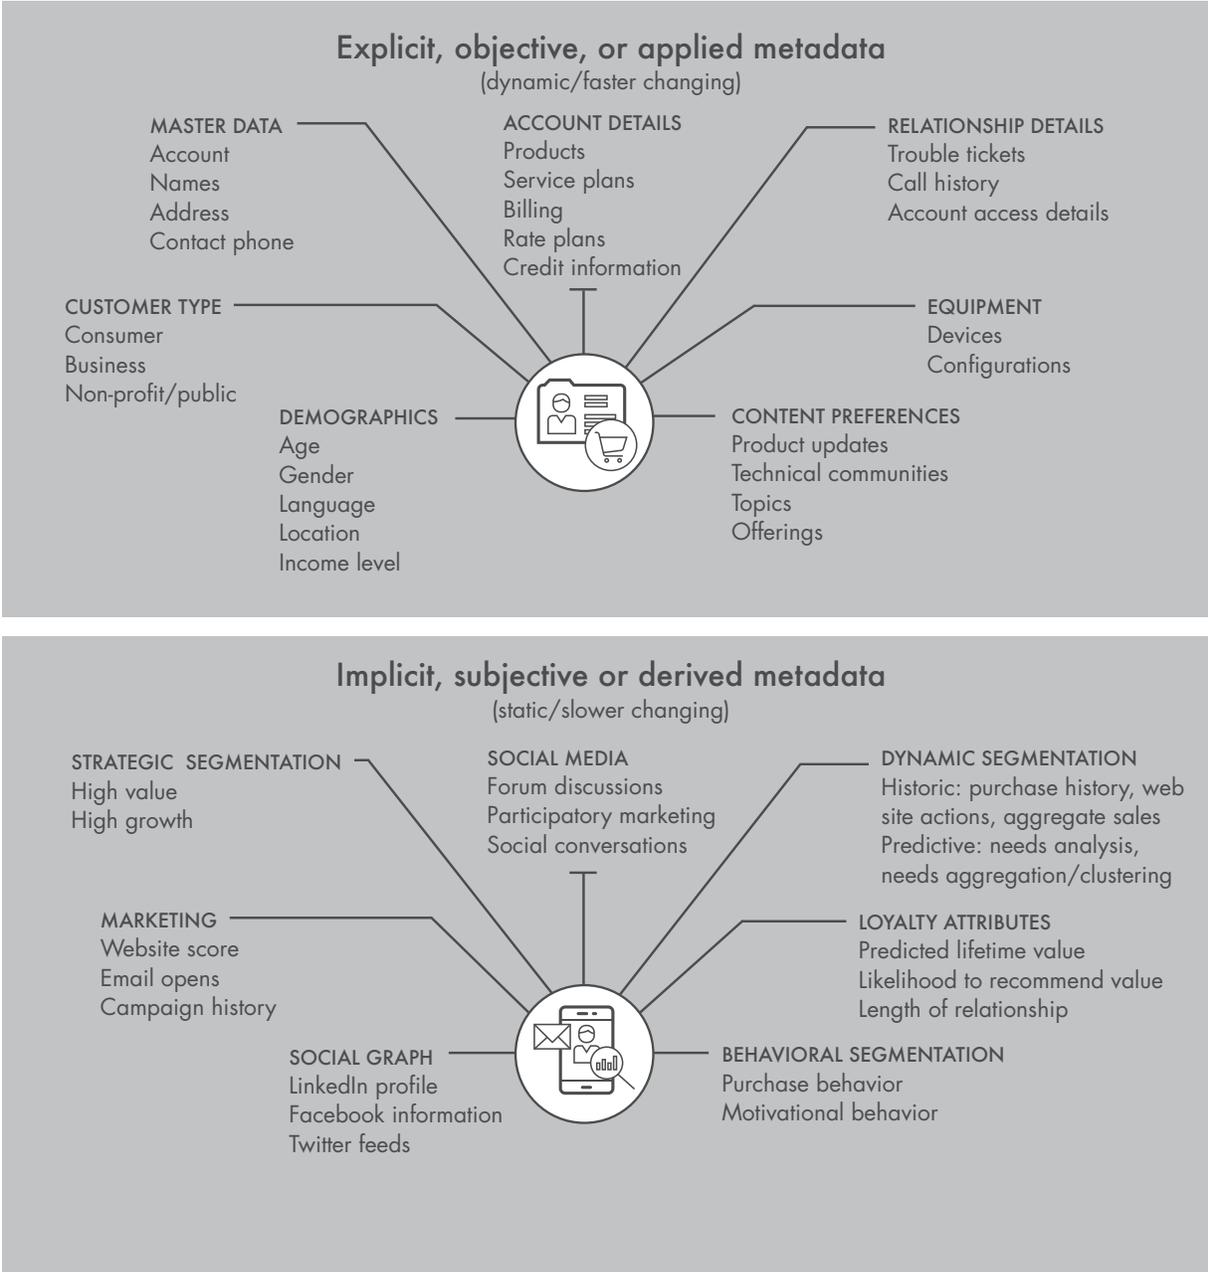
By recruiting test subjects for each persona, you can determine which elements work for all customers and which ones need to be adapted for subsets of customers. Testing across representative audiences tunes design, informs a personalization framework, and is the foundation for metrics-driven governance.

### **Audience Attributes**

Tracking audience attributes allows you to move from design to site personalization. Those attributes can become quite complex. You can describe a customer based on details of products or services they have purchased in the past, their history with customer service, the types of equipment they may own, their status as businesses or consumers, their demographic details, and the types of content they are interested in.

As we mentioned in chapter 3, you can infer other characteristics that are less clearly defined, that may be rapidly changing, and that arise from a combination of human judgment and AI. For example, loyalty may be a composite measure based on length of relationship and past behavior. Or the customer may be described as “high value” based on frequency or volume of purchases. Figure 5-1 shows how customer metadata can reflect both explicit and implicit characteristics.

Figure 5-1: Explicit and Implicit Customer Data



**Customer Data Platforms and Customer Information Flows**

Because so many different systems touch users, sites use a crucial piece of software that constantly takes in the “data exhaust” from other systems: the customer data platform (CDP). This software can make decisions about how to customize a site or aggregate signals for input to AI tools. In order for the CDP to function, the different systems have to consistently define attributes of the customer and their interaction using identical or analogous terminology.

The definitions of customer characteristics constitute an attribute model. Since each system has some representation of what customers are doing—whether clicking banner ads or opening email and searching on websites—these systems must use the same names for the same concepts. However, these tools typically come from different companies, so they inevitably will use different names for the customer attributes. For example, one system might track location as “geography” while another calls it “region.” The model has to be defined to account for all customer data and how it flows among systems.

Some systems act on and change customer data—for example, when the customer orders a product. That new information goes into multiple systems, including the customer relationship management (CRM) system and the enterprise resource planning (ERP) system. A CDP might take that into account and then direct the content management system (CMS) to provide content more appropriate for a current customer rather than for prospects. The CDP may also cause the campaign management system to acknowledge that the target is a current customer. For example, if a customer responded to a campaign and filled out a form to download a white paper, the CDP could use that form to add profile elements such as “topics of interest.” The campaign management tool would then communicate back to the CRM system, adding the topics of interest to the customer’s record. The CDP and other systems would now have that data available to further inform future contact with the customer. Mapping information flow among technologies is a critical element of CDP deployment. For more detail, see the section on ontologies in ecommerce later in this chapter.

## THE ESSENTIALS OF SITE CUSTOMIZATION

So far in this chapter, I’ve discussed the product and customer data that organizations collect and how sites customize search and navigation based on that data. But customization goes beyond navigation. So let’s take a look at how sites will flex based on who’s visiting and what they’re looking for.

Consider how your organization reacts to the signals that indicate which kind of customer is visiting. For some organizations, data signals generate millions of customized audiences and offers, each one specifically tailored to exactly what that customer needs. Sites like wayfair.com use this mass personalization approach to optimize each user's experience on the site.

### **Designing for Automation: Begin with Artisan Processes**

Automation is not magic. It starts with a complete understanding of the path humans follow and the ways a site can react.

When you set out to automate a personalized ecommerce experience, no system will know what a particular user or audience needs without the crucial factor of human insight. A marketing specialist who knows their customer will have to decide what message or part of a message they *think* will resonate—and then test it. They can then handcraft the message and try a variation, just as an artisan uses knowledge of their craft to create something that will engage with another human. The marketer will then try other variations and learn what other items might work and which ones do not.

Once you understand what drives messaging and design unique engagement components for it, you can experiment. Here is where the metrics and feedback allow for optimization—*handcrafted optimization*. This is not practical or scalable. But it is the beginning of understanding how to vary the design elements that can be recombined by a system and process and, ultimately, an algorithm.

One organization had 20 components that could be varied on a landing page: different images, offers, calls to action, phrasing, and headlines that could be rearranged and reassembled like Lego blocks. Suppose these elements could be arranged in five slots on the page, with four possible elements in each slot. The number of possible combinations of elements that could make up that page is four to the fifth power ( $4 \times 4 \times 4 \times 4 \times 4$ ) or 1,024—over 1,000 possible variations of that page from those Lego blocks.

After breaking the messages up into reusable building blocks, the blocks need to be structured and organized through—as you probably guessed—the ontology. The ontology has the instructions for how the pieces fit together by virtue of its models of the content. Think of how tiles on a

mosaic floor might fit together. The pieces have regular patterns that allow them to fit together within predetermined templates that have slots for each of the pieces.

What's the best way to assemble these hundreds or thousands of possible combinations? This is where machine learning algorithms come in. The input is the messaging architecture and a collection of components for each design element. The output is a random assembly of those components, which is then tested with audiences. The algorithm analyzes the data, picks the winning elements for that audience and journey stage, and then continues to experiment across other audiences, selecting winners in a Darwinian process. Over time, the system can generate hundreds or thousands of personalized experiences. Other algorithms can determine which products are presented once the customer clicks through on an offering. What that customer sees will depend on past purchase analysis, behaviors of similar audiences, and recent activity signals. You can do the same testing with ad copy, promotional offers, and sequences of messages.

This analysis is based on large-scale experimentation. It uses a form of AI called a *neural network*—a computing structure that “learns” patterns by observing the results of experiments. The neural network tests various weights to components and audience attributes until they produce optimal outputs. The network has a target of maximizing revenue and will experiment until it reaches the best performance possible with those components.

### **A Detailed Architecture Allows Customization at Scale**

Personalization at this scale is a process of continually testing and recombining elements of design, messaging, and offerings. Marketers, of course, lack the bandwidth to customize messages across hundreds or thousands of audiences. What they need to develop is a messaging *architecture* that the AI can optimize across these audiences.

Even if the organization lacks the processes to develop a messaging and component architecture, there are many reasons to design potentially personalized content based on a large number of descriptors. If segment data is detailed, you can perform analytics to identify underserved markets and customers. As capabilities evolve, fine-tuning

offers with machine learning will become a standard practice. Other tools in the organization's technology ecosystem will benefit from fine-grained audience definitions. New ways of interacting with customers will evolve quickly.

This level of detail also lends itself to revealing new insights about how the customer responds to offers and messaging. If you lump a large number of customers into one big bucket of "new customers," it will not be possible to understand their differences across other characteristics and descriptors. Remember that, ultimately, humans will not be selecting attributes and personalizing content based on them—algorithms will. A customer attribute model defines the data features that the machine learning algorithm will use. Machine learning algorithms perform best with a lot of variables; the more customer features we define up front, the better the algorithm can perform.

AI can actually help identify audiences with common responses to offers. Machine learning can look at all of the signals from customer characteristics and seek latent or hidden attributes that explain how offer characteristics relate to audience responses. The algorithms can predict who else might respond similarly to the same characteristics or identify new audiences based on subtle common details in prior audiences. Those esoteric customer features could depend on complex relationships with dozens of variables, but they indicate that two audiences are similar. The algorithm can identify these "look-alike" audiences through combinations of data gleaned from shopping patterns, demographics, industry, job role, interests, purchase history, and other user signals.

Initially your organization will not act on all of the signals. You can simplify your approach by customizing a broad model of the domain for a particular use case and application; the added granularity in the data will improve future flexibility and adaptability. The ultimate goal is for the system to read the customer's digital "body language," just as a star salesperson might use their knowledge of all the customer's past history and tendencies and behavior in the moment to adjust exactly how to best communicate with that customer.

## ONTOLOGY FOR ECOMMERCE

If data is the foundation of a customized site experience, wrangling that data from multiple systems is increasingly the most difficult ecommerce challenge. The solution to that challenge lies with the one corporate element that ties together all the data: the ontology.

Each of the data elements that describe the customer needs to be harmonized with the other customer experience and engagement systems. To function together, the systems require common ways of describing these elements and organizing principles. The ontology is where these elements are synchronized.

Machine learning programs refer to these common elements as “features.” Though the machine learning program figures out some of these categories and classes itself, those features are not sufficient. Defining the features that are unique to your organization is key to your differentiation and competitive advantage.

Product categories are part of the ontology, as are the various vocabularies that define features and specifications. Those relationships are often contained in specialized ontology management tools, or can be embedded in master data, metadata, and CRM systems. The important piece is to be consistent when possible, and to differentiate where necessary, depending on the application.

For example, as we’ve described earlier, product classifications have to be consistent across the ecommerce system, the CRM system, the content management system, and the enterprise resource planning (ERP) system. This does not mean they are exactly the same in each of those systems. The ERP system typically uses a structure that lends itself to financial reporting, while the CRM system may use a more territory-specific ordering or subject matter expert grouping. Merchandisers need a view of product collections that is different from the one that customers see, and product classifications for the merchandiser can provide that. Those variations need to be intentionally designed rather than left to independent groups making decisions as they develop, configure, or deploy their tools—and the ontology needs to rationalize and connect them.

What about the product information management (PIM) system? It’s

not a substitute for the ontology. While a PIM will contain many features (in the form of metadata and product models), the ontology is what translates the important elements across systems and provides knowledge relationships that are unlikely to be contained in the PIM. The PIM becomes more powerful because of ontological elements.

### **Ecommerce Depends on a Robust Collection of Attributes in the Ontology**

The ontology might include products, product categories, campaigns, market segments, customer types, content types, channels, and suppliers. The ontology contains language variations through “equivalence relationships”—terms that represent the same concept but that may not be actual synonyms. For example, “opacity” and “transparency” could be equivalence terms because even though they have different meanings, they represent the same concept.

Other equivalence terms include terminology for use by the search engine: technical terms, acronyms, abbreviations, slang terms, and other non-preferred terminology. Ontologies contain references to resources that the user might not think of; they are like reference librarians that know what to suggest when the user enters an ambiguous query. They can also contain cross-sell, upsell, and other related product relationships, inherited from a PIM system or separately.

The ontology contains customer roles and their associated industries, interests, and problems. This information allows the search engine to return results in the correct context. Knowing the customer’s industry, for example, allows for disambiguation of “mold stripping” in the example we described earlier.

Content is tagged with metadata from the ontology, which should be constructed to be the source of truth for all organizing principles in the enterprise. Product attributes are sourced from the ontology as well, and are ultimately applied to product records in the PIM system.

AI and machine learning technologies can aid in the process of applying the correct metadata to products by recognizing patterns and clustering and similar products so that an analyst can classify them. Text analytics machine learning programs can extract product data from source documents like

PDFs and engineering specifications. The ontology contains the reference data and gives the machine learning technology the correct lists from which to recognize attribute data. Conversely, these same tools can also begin to extract candidate terms and attributes from text to start or enrich the ontology. Figure 5-2 shows the start of a sample ontology.

Figure 5-2: Sample Elements of an Ecommerce Ontology

## Ontology Development

### RELATIONSHIPS:

- Products for processes
- Tasks for customer types
- Solutions for processes
- Solutions for industries
- Interests for customer types

COMPETITORS	ABC (Organization)		
<ul style="list-style-type: none"> <li>• Grainger</li> <li>• Wolseley</li> <li>• Fastenal</li> <li>• ...</li> </ul>	<ul style="list-style-type: none"> <li>• Tech support</li> <li>• Merchandizing</li> <li>• Direct sales</li> <li>• ...</li> </ul>		
PRODUCT CATEGORIES	CUSTOMER TYPES	INTERESTS	INDUSTRIES
<ul style="list-style-type: none"> <li>• Abrasives</li> <li>• Clamping</li> <li>• Fasteners</li> <li>• Fleet maintenance</li> </ul>	<ul style="list-style-type: none"> <li>• Procurement manager</li> <li>• Engineer</li> <li>• Maintenance manager</li> <li>• ...</li> </ul>	<ul style="list-style-type: none"> <li>• Prototyping</li> <li>• MRO</li> <li>• Re-stock</li> <li>• ...</li> </ul>	<ul style="list-style-type: none"> <li>• Government</li> <li>• Manufacturing</li> <li>• Healthcare</li> <li>• ...</li> </ul>
PROCESSES	SOLUTIONS	TASKS	...
<b>Attaching</b> <ul style="list-style-type: none"> <li>• Adjoin</li> <li>• Adhere</li> <li>• Affix</li> <li>• Anchor</li> <li>• ...</li> </ul>	<b>Solution Categories</b> <ul style="list-style-type: none"> <li>• Inventory management</li> <li>• Metalworking</li> <li>• ...</li> </ul>	<b>Tasks</b> <ul style="list-style-type: none"> <li>• Cutting titanium</li> <li>• Mold stripping</li> <li>• ...</li> </ul>	<ul style="list-style-type: none"> <li>• ...</li> <li>• ...</li> <li>• ...</li> </ul>

## SPECIAL CONSIDERATIONS IN B2B COMPANIES

In business-to-business companies, especially, ecommerce is about more than product sales. Customers are choosing and purchasing services of all types; even in complex sales that typically require guidance and knowledge, they are conducting research online in advance of the sale. This is true for industries that have historically sold through personal relationships, including those in which sales have previously been dependent on human judgment and expertise.

In earlier times, paper catalogs contained product clues, or cues, that

customers depended on. These wayfinding cues are different in the online world, because digital processes drastically increase the need for highly detailed information structures. When product data is missing, products drop from search results or do not display correctly in online comparisons. In other words, having the wrong product information design makes your products invisible.

Whether you are a retailer, a B2B enterprise, or a distributor, your ecommerce system performs a critical function in your customer journey by helping customers find and learn about your products. And if they can't find it, they can't buy it.

B2B ecommerce is less mature than B2C due to the complexity of the sales, the nature of large product catalogs, the unique needs of customers, and the traditional reliance on selling through human relationships. Industrial markets can be quite diverse and complex, which makes fine-tuning the online user experience more challenging. However, since B2B customers are also B2C customers in their non-work roles, they have become accustomed to high-quality web experiences. As websites have evolved, user expectations have increased. Digital natives—who have lived their entire lives with technology—expect to do everything through smartphones and to have easy access and “Google-like” search experiences.

As a result, B2B customers in all sectors are increasingly expecting a seamless, consistent experience. Organizations can compete based on how well they understand the types of problems their customers need to solve and how to best organize and present that information. Success in creating such an experience requires product taxonomies and data structured in a way that is intuitive to the customer.

The next chapter contains much more detail on the sales funnel, especially as it applies to B2B companies.

## OVERLAPPING TOOLS AND ORGANIZATIONAL MATURITY

Complex ecommerce functionality includes elements of product information management, digital asset management, customer information management, content management, marketing and promotions management,

email management, social media marketing, shopping cart functionality, payment and order management integration, data quality and governance, and more. Personalization mechanisms can and do span most of these functions. Many organizations have chosen a modular approach to building their platforms, with many best-of-breed applications being integrated with web services. This modular approach can optimize functionality with specialized approaches to niche problems, but it increases cost and complexity.

Ecommerce technology platforms can include a range of functions even as they integrate with other tools that provide specialized functionality. For example, some ecommerce tools include product information management (PIM) technology, but others leave that functionality to PIM vendors. Some PIM tools include digital asset management (to manage images, video, and other rich media). Specialized and customized tools make the environment more complex, adding cost and reducing agility.

How pervasive are these overlaps in function? Consider an exercise I conduct in workshops on information management strategy. I lead participants through the process of creating a content model for an invoice to be accessed in a content management environment. As I encourage participants (“What’s missing? What else?”), they add details and more and more metadata. Then they suddenly realize that they have left the confines of content management and are actually designing an accounting system, not just a document structure. Details around line items and quantities, shipping addresses, pricing, and the like do not belong in the content management environment; they belong in the invoicing application.

Then they begin to ask themselves: What data did the customer actually need? What functionality? What tasks and problems are they trying to solve? And what is the most appropriate tool for the job? It is possible to design more detail in the PDF of the invoice, but why not just retrieve it from the accounting system?

This exercise illustrates the importance of “fitness to purpose”: designing and configuring applications to have appropriate functionality rather than shoehorning in things that the package or platform was not designed for. You must choose technology in the full context of what is in place today and what will be needed in the future.

You should expect some overlap in the functionality of the various ecommerce systems, such as an ecommerce suite that includes some content management but that might also integrate with a content management system. The key is to define scenarios and use cases and to determine where those scenarios can be best handled with the greatest flexibility and least amount of complexity. Adding new code adds cost and complexity. While it is possible to build accounting functionality in a CMS, it would be costly and confusing to do so, since that is far from core functionality.

An ontology manages, translates, and integrates the inevitable design differences among customers, content, and product attributes. It helps speed development and integration by providing the reference terminology for each of the components of an ecommerce suite.

### **Maturity Assessment**

One useful tool in evaluating ecommerce capability is the maturity model. Since ecommerce requires orchestration of information across multiple systems and processes, you must consider each of the major systems in relation to the critical functions that comprise the ecommerce technology stack. Orchestration is more than integration. It is the finely tuned digital machinery that manufactures the digital relationships with your customers.

The critical ecommerce functions include the following:

- **Product launch and item onboarding.** How long does it take to bring a new product online? Is this a fast, seamless automated process or one that takes many weeks and manual activities to accomplish?
- **Product information management.** How detailed is the product data model? Does it include the attributes that are important to target customers? Are solution and merchandising attributes and relationships defined? Can this data be sourced or distributed in a syndicated way? Specifically:
  - Are multiple taxonomies in use?
  - Do multiple taxonomies lead to an inconsistent user experience? (Different parts of a site sometimes use different taxonomies.)

- Does the information architecture use industry best practices?
- Does it leverage industry standards?
- Is content managed in a CMS?
- Is there a PIM system?
- Is the information architecture consistent across channels and devices?
- **Product configuration.** Is product information static, or does it allow for configuration and customization? Are there complex rules engines with configuration constraints? Does the system use an expert advisor bot for complex, custom-built-to-order products?
- **Analytics.** What is the level of sophistication of analysis of customers, segments, traffic, conversions, multivariant testing, and personalization metrics?
- **Commerce and order management.** What capabilities are available to customers? Does the system include conversational advisor order management, shopping-cart-level promotions, and segment personalization attribute triggers?
- **Payments and pricing.** Some less mature sites require users to make a call to obtain pricing information and do not have commerce capability. Others require registration or manual setup with a salesperson. More advanced capabilities include purchase workflow/approval, one-click ordering, automated replenishment with negotiated pricing, multidimensional personalization, and a seamless omnichannel experience.
- **Data quality.** For example, a PIM system can report on how completely product data is filled out using visualization tools, audit trails, data quality services, and compliance mechanisms. Ideally, this would link to business outcomes and automated exception reporting.

Table 5-1 describes maturity levels across each of these dimensions.

Maturity will likely be uneven, with some departments and processes being more advanced than others. It may not be necessary to fully optimize each activity. The process requirements, website capabilities, and engagement strategy may not require Stage 5 maturity.

Consider that the time required for a large enterprise to go from one maturity level to another is typically six months to one year. If the organization is at Stage 1 and the process requires Stage 4, but the timeline and budget requires attaining that maturity in six months, you will need additional resources. If budgets don't allow for that acceleration, you need to scale down the project.

### **Maturity Checklist for Ecommerce**

To help you operationalize the maturity rubric I just described, Table 5-2 presents a useful checklist of ecommerce elements. The list is not meant to be exhaustive, but you can use it to place some functions along a continuum of complexity and maturity, with the staffer most likely to be able to help you assess maturity shown in parentheses. Create your own model based on the important elements, and grade where your current processes are versus the desired future state. That will provide some sense of the gap between the current state and the future.

Table 5-1: Example of B2B Ecommerce Maturity Levels

Capability \ Maturity	1-Unpredictable	2-Aware
Product Launch and Item Onboarding	Cycle times in weeks or months; manual processes	Multiple out-of-synch systems need to be updated and reconciled
Product Information Management	Poor data quality; manual validation; limited transformation	Category hierarchies and attributes identified; semi-automated data quality processes
Product Configuration	Basic static product selection without configuration	Base variant configuration
Analytics	No analytics on customer purchases	Basic customer segmentation and purchase analytics
Commerce and Order Management	No transaction capability	Basic commerce with persistent shopping cart; abandoned cart triggered promotions
Payments and Pricing	Call for pricing; no public pricing; no commerce capability	Pricing available with registration; manual setup with salesperson; basic catalog site
Data Quality	No data quality functionality	Limited data quality reporting

3-Competent	4-Synchronized	5-Choreographed
Multiple content streams triggered by product launch	Tuned taxonomies for channels; A/B testing for launches	Seamless cross-channel product launch; cycle times in days
Business rules built into product data quality; manual hierarchy mapping to channels	Data quality feedback to vendors; consistent cross-channel mapping	Many-to-many vendor-to-channel cross-mapping; integration with domain analytics
Base variant and product customization capability	Complex configuration with rules engine and configuration constraints; business and design tool integration	Complex built-to-order integrator expert advisor bot; nested configurations; design and business tool integration
Reporting on-page traffic; popular products; conversion rates; orders and search terms	A/B testing with feedback adjustments; dashboards for UI; merchandizing; campaigns and product bundle performance testing	End-to-end KPIs linked to multi-variant testing and remediation; personalization optimization metrics
Cross-sell and upsell merchant configurable with shopping basket recommendations	Cross touchpoint promotion management; search-based merchandizing	Conversational advisor order management, including shopping cart level promotions; segment personalization attribute triggers
Basic personalization; purchase workflow/ approval; one-click ordering; bulk pricing	Automated replenishment with negotiated pricing; multi-currency; multi-language	Multidimensional personalization; seamless omnichannel experience
Integrated data quality metrics via PIM or PIM integration; attribute fill; completeness	Visualization tools; audit trails; integration with third-party data quality services; compliance mechanisms	Linkage to business outcomes; automated exception reporting; queue management and workflow; semi-automated fill

Table 5-2: Maturity Checklist for Ecommerce

<b>Supplier Onboarding</b>	<ul style="list-style-type: none"> <li>• Have suppliers provided an inventory of available metadata? (sourcing manager)</li> <li>• Have product data models been validated against merchandizing requirements? (catalog manager, merchandiser)</li> <li>• Are data quality review processes and scorecards in place? (procurement, data quality manager)</li> <li>• Are supplier data curation responsibilities defined in a Master Services Agreement? (procurement)</li> <li>• Are data quality performance standards linked to procurement performance measures? (procurement)</li> </ul>
<b>Product Data Onboarding</b>	<ul style="list-style-type: none"> <li>• Are product managers getting needed content and data? (product data operations)</li> <li>• Is product data complete and consistent? (data quality manager)</li> <li>• Is product data appropriate for merchandising? (merchandiser)</li> <li>• Do product attributes help users make selections? (information architect)</li> <li>• Are technical specifications organized with product content? (product data operations)</li> <li>• Are images and rich media assets tagged with product data? (catalog manager)</li> </ul>
<b>Demand Generation Metrics</b>	<ul style="list-style-type: none"> <li>• Are social media campaigns assigned tracking codes? (marketing analyst, marketing manager)</li> <li>• Are promotions measured against baselines? (marketing analyst, marketing manager)</li> <li>• Are email engagement metrics monitored and acted upon by product marketers? (marketing analyst, marketing manager)</li> <li>• Are campaign effectiveness metrics benchmarked across categories? (marketing analyst, category manager)</li> <li>• Are engagement metrics aligned with customer lifecycle? (marketing analyst, category manager, information architect)</li> </ul>

Table 5-2 (Continued)

<p><b>Omnichannel Experience</b></p>	<ul style="list-style-type: none"> <li>• Are store promotions consistent with online promotions? (marketing analyst, marketing manager, regional manager)</li> <li>• Can customers find products in stock at stores? (marketing analyst, regional manager)</li> <li>• Can customers select products on mobile devices? (information architect)</li> <li>• Can customers begin transaction on device and complete in store? (merchandiser, ecommerce manager)</li> <li>• Are online preferences retained for personalized shopping? (ecommerce director, information architect)</li> </ul>
<p><b>Self-Service Metrics</b></p>	<ul style="list-style-type: none"> <li>• Is self-service content call deflection measured? (call center manager, knowledge base owner)</li> <li>• Is engineer knowledge base contribution measured? (knowledge base owner)</li> <li>• Is support content quality measured? (support manager, knowledge base owner)</li> <li>• Is user-generated content quality measured? (community manager, content manager)</li> </ul>
<p><b>Content Operations</b></p>	<ul style="list-style-type: none"> <li>• Is editorial content refresh workflow defined? (content manager)</li> <li>• Do expiration dates drive action? (editorial director, content manager)</li> <li>• Are asset rights managed and respected? (content manager)</li> <li>• Are current promotions lifecycles automated? (marketing manager, merchandiser)</li> <li>• Are category manager objectives communicated to content operations? (category manager, editorial director)</li> <li>• Is support content presented at point of need? (content manager, knowledge base owner)</li> <li>• Are user-generated content ingestion workflows defined? (community manager, content manager)</li> <li>• Is content tagging and association automated? (content manager, ecommerce director, merchandiser)</li> </ul>

Table 5-2 (Continued)

<b>Personalization Strategy</b>	<ul style="list-style-type: none"> <li>• Are buyer needs documented and user personas defined? (merchandiser, marketing manager, information architect)</li> <li>• Are tasks and objectives aligned with specific content? (content manager, information architect)</li> <li>• Are related content aggregation rules defined? (editorial manager, content manager, information architect, marketing analyst, merchandiser)</li> <li>• Are personalization rules developed for buyer segments? (information architect, marketing manager, marketing analyst, merchandiser)</li> <li>• Is content dynamically assembled based on user behavior? (information architect, marketing manager, marketing analyst, merchandiser)</li> </ul>
<b>Digital Assets</b>	<ul style="list-style-type: none"> <li>• Are spec sheets associated with product content? (product data operations)</li> <li>• Are engineering drawings available when appropriate? (merchandiser)</li> <li>• Are assets organized for retrieval and reuse? (information architect)</li> <li>• Are generic versus specific image requirements defined? (merchandiser)</li> <li>• Are marketing brochures organized appropriately? (information architect, marketing manager)</li> <li>• Are high-resolution images available as appropriate? (information architect, content manager)</li> </ul>
<b>Analytics</b>	<ul style="list-style-type: none"> <li>• Are content performance metrics monitored and embedded in governance processes? (information architect, marketing manager, marketing analyst, merchandiser)</li> <li>• Are user departure paths understood? (information architect, marketing analyst)</li> <li>• Is search effectiveness measured? (information architect, content manager)</li> <li>• Are conversion and purchase behaviors monitored across dimensions and factors? (marketing manager, marketing analyst, merchandiser)</li> </ul>

Table 5-2 (Continued)

<b>Content Architecture</b>	<ul style="list-style-type: none"> <li>• Are metadata models fully defined? (information architect, content manager)</li> <li>• Are controlled vocabularies harmonized? (information architect, content manager)</li> <li>• Does content architecture support customer experience—that is, do taxonomies and content models support customer tasks? (information architect, marketing manager, marketing analyst, merchandiser)</li> <li>• Are content and product data models and relationships aligned? (information architect, PIM system manager)</li> <li>• Does the content architecture support a dynamically generated experience: (information architect, marketing manager, marketing analyst, merchandiser)</li> <li>• Are all cross-sell opportunities supported with rich content? (merchandiser, content manager, editorial director)</li> </ul>
-----------------------------	--

The job of any ecommerce site is, ultimately, to sell. But sales requires specialized tools, as I'll describe in more detail in the next chapter.

## TAKEAWAYS FROM CHAPTER 5

In this chapter, I've shown how AI has its broadest and most powerful application in ecommerce sites. But that's also where the product and customer data, design, and planning can go off the rails without appropriate attention to the ontology. These are the main points in this chapter:

- Ecommerce—even when it's not where the sale gets made—is where ontology-powered AI can have its biggest impact.
- In a modern site, search and site navigation have become a customized, dynamically updating conversation.
- Data is the foundation of modern ecommerce, and the quality of the experience depends on the quality of the data.

- Product taxonomies that reflect not just product features, but relationships, make sophisticated ecommerce possible.
- Standardized product taxonomies enable interoperability with the supply chain, but differentiated taxonomies can create competitive advantage.
- Customer classification begins with personas and mental models.
- Testing must include all types of potential customers and personas.
- Customer data platforms assemble all relevant customer data from multiple systems.
- Site customization starts with artisanal experiments but continues with improvements tested by neural networks.
- Ecommerce depends on a robust collection of attributes in the product ontology.
- Customer and product information systems inevitably overlap; efficiency requires straightening out their relationships with each other.
- You should assess the maturity of your customer systems to determine where improvements will have the biggest impact.



## GLOSSARY

**“ABOUT-NESS.”** Characteristic that describes the classifications and qualities of data elements in an ontology; metadata “about” content, documents or any source or piece of information. *See also* “is-ness.”

**ACCIDENTAL ONTOLOGIES.** Data architectures that arise naturally from individual technology projects, but without an overarching organizational principle.

**ARTIFICIAL INTELLIGENCE (AI).** A broad term representing multiple classes of algorithm that identify patterns in data and emulate some aspects of human judgment to draw non-obvious conclusions from them.

**ARTIFICIAL NEURAL NETWORKS (ANNS).** *See* neural networks.

**ATTRIBUTE MODEL.** The set of descriptors (metadata) that describe content, products, customers, knowledge, or structured data sources. *See also* customer attribute model.

**AUDIENCE.** A collection of users (for example, customers or employees) attempting to solve a problem (for example, finding a product to buy or verifying a policy).

**AUGMENTED REALITY.** A technology that allows a visual display to overlay digital information on a view of the real world.

**BANTS CRITERIA.** An acronym for the five criteria used to qualify leads: budget, authority, need, timeline, and strategic fit.

**BIG DATA.** Massive collections of information that can't be analyzed with traditional methods due to demanding levels of volume, variety, and velocity.

**BOTS.** Automated tools that perform activities or interface with customers or employees. *See also* chatbots.

**CAD/CAM (COMPUTER AIDED DESIGN/COMPUTER AIDED MANUFACTURING).** A set of tools that enables the virtual creation, visualization, and manipulation of objects to be manufactured.

**CATEGORIZATION.** Classification applied to similar types of data, such as product descriptions or documents.

**CDP.** *See* customer data platform.

**CHATBOTS.** A class of bots that can conduct conversations with customers or employees.

**CHIEF DATA OFFICER (CDO).** A C-suite executive whose responsibility is to improve the consistency and quality of data in an enterprise.

**CLUSTERING.** Identifying similar content or documents—for example, to help contextualize search queries.

**COBOT (COLLABORATIVE ROBOT).** A robot that provides assistance working alongside a human.

**COGNITIVE LOAD.** The amount of mental effort that a customer or employee

needs to invest to accomplish a task. Interfaces that pay too little attention to user experience often end up generating an increase in cognitive load.

**CONFIGURE-PRICE-QUOTE (CPQ) SYSTEM.** A tool that allows selection of product options to meet user needs and preferences, and that can streamline the selection and quoting process by recommending and pricing appropriate combinations of features.

**CONTENT MANAGEMENT SYSTEM (CMS).** A system for managing, assembling, and distributing content on a website or other customer engagement application.

**CONTENT MODEL.** A structured representation of concepts within a document, defined as metadata regarding “is-ness” and “about-ness”.

**CONTENT MARKETING.** A marketing practice that creates useful content for customers to attract and retain their attention.

**CONTEXT.** The environment in which data or content is considered or applied; to operate on the data, AI algorithms often need to know the context to operate on the data. Users need content in the context of their goal.

**CONTEXTUALIZATION.** The process of identifying the context in which data is considered or applied.

**CONVERSATIONAL MARKETING.** A practice in which automated tools use natural language interactions to converse with customers.

**CRM.** *See* customer relationship management system.

**CROWDSOURCING.** A process that uses shared comments in an online social network to gather information and ideas for a project.

**CUSTOMER ATTRIBUTE MODEL.** A set of descriptors that represent the customer’s interests, needs, tasks, objectives, role, history, propensity to buy, and so on;

the “is-ness” and “about-ness” of a customer described in metadata terms.

**CUSTOMER DATA PLATFORM (CDP).** A system that aggregates customer data (the customer’s “digital body language”) from multiple sources for the purpose of improving engagement.

**CUSTOMER EXPERIENCE (CX).** The sum total of all experiences the customer has in their interaction with an organization—both online and offline; the discipline of analyzing such experiences.

**CUSTOMER JOURNEY.** The path that a customer takes as they interact with the company. *See also* high-fidelity customer journey, journey map.

**CUSTOMER LIFECYCLE.** The set of broad stages a customer goes through as they interact with your organization—for example, learn, choose, purchase, use, maintain, and recommend.

**CUSTOMER MODEL.** A description of the customer in terms of everything the organization knows about them. *See also* customer attribute model.

**CUSTOMER RELATIONSHIP MANAGEMENT (CRM) SYSTEM.** A system that holds information about customers and their interactions with an organization.

**DARWIN INFORMATION TYPING ARCHITECTURE (DITA).** A document component architecture that makes information more easily consumable by computers.

**DATA DICTIONARY.** A description of the contents, format, structure, and relationships contained in a database.

**DATA HYGIENE.** The discipline of keeping data accurate and up to date to avoid errors, costs, and inefficiencies.

**DATA MODEL.** A broad term used to represent the makeup of structured, unstructured, and semi-structured information throughout the enterprise.

**DATA QUALITY.** A measurement of the trustworthiness, completeness, accuracy, and timeliness of a set of data.

**DIGITAL ASSET MANAGEMENT (DAM) SYSTEM.** A software system used to manage rich media digital content (such as images, video, drawings, and photos) for deployment in marketing communications, for example.

**DIGITAL TRANSFORMATION.** A corporate initiative to reengineer a company's processes and culture to be fully digital and integrated.

**DIGITAL TWIN.** A virtual representation of a physical product that can be used in tests and simulations.

**DISPLAY TAXONOMY.** A product classification design used to aid site navigation. *See also* navigational hierarchy.

**DOMAIN MODEL.** A structured view of data and content in a given information area (such as customers, vendors, products, or finance) or process (such as product onboarding, marketing campaign development, or customer service).

**EMPLOYEE JOURNEY.** The steps an employee takes to solve a problem; analogous to customer journeys.

**ENGAGEMENT STRATEGY.** How your organization differentiates from the competition through messaging, channels used to reach its audience, how it delivers value, and the character of the customer relationship.

**ENTERPRISE RESOURCE PLANNING (ERP) SYSTEM.** A software system that companies use to manage resources and integrate other systems, such as accounting, inventory management, and procurement systems.

**ENTITY EXTRACTION.** An AI technique that identifies concepts and associated terms in text.

**EQUIVALENCE TERMS.** Different terms referring to the same concept in the ontology.

**EXPLICIT (CUSTOMER) DESCRIPTORS.** Facts about customers stored in customer systems. *See also* implicit (customer) descriptors.

**EXPLICIT KNOWLEDGE.** Knowledge documented in knowledge bases.

**FACETS.** Options presented in a search that refine search queries for faster resolution. Facets use product and content attribute models.

**FEATURES.** In the context of customer attribute models, a feature is a type of customer or product descriptor. Also described as attributes or metadata. In the context of text analytics, it refers to ways of describing the information in a document. In machine learning, a feature is a property or characteristic in the data. Feature engineering is the selection of characteristics that have a meaningful impact on the output of the algorithm.

**FRICTION.** Anything that impedes the efficient functioning of an enterprise.

**GROUPWARE.** A category of software, now largely obsolete, that was designed to improve information sharing. Example: Lotus Notes.

**HELPER BOT.** A bot designed to aid employees—for example, call center workers.

**HEURISTICS.** Rules of thumb, standards, and rubrics for evaluation—for example, in the context of a maturity assessment, user experience, data structures, or content organizing principles

**HIERARCHY.** A classification system, such as a taxonomy, consisting of parent–child or whole–part relationships. *See also* navigational hierarchy.

**HIGH-FIDELITY CUSTOMER JOURNEY.** A representation of the customer’s needs in data terms, including attributes and descriptors that indicate their role, buying state, interests, demographics, goals, and even state of mind.

**HIGH-FIDELITY (CUSTOMER) JOURNEY MAP.** A validated model of the high-fidelity customer journey featuring data elements about the customer at every stage.

**IMPLICIT (CUSTOMER) DESCRIPTORS.** Inferences about customers based on other data. *See also* explicit (customer) descriptors.

**INFERENCE ENGINE.** A mechanism that can answer questions beyond those originally programmed into it.

**INFORMATION ARCHITECTURE (IA).** The structural design of information that comprises the user experience or that enables digital capabilities.

**INFORMATION METABOLISM.** An analogy used to describe the flows of information within businesses.

**INFRASTRUCTURE PROGRAMS.** Programs that improve the basic functioning of technology within an enterprise but that do not have a direct line of sight to the customer.

**INTELLIGENT VIRTUAL ASSISTANT.** *See* virtual assistant.

**INTENT.** The task a user of a system is trying to accomplish, typically identified from a finite list of possible desires.

**INTERNET OF THINGS (IOT).** The broader computing environment that results from the proliferation of connected devices.

**"IS-NESS."** Characteristic that describes the fundamental nature of data elements in a content model and the ontology. *See also* "about-ness."

**JOURNEY MAP.** A high-level description of the customer journey. Also called a customer journey map. *See also* high-fidelity journey map.

**KNOWLEDGE ARCHITECTURE.** An ontology and associated content models for organizing knowledge.

**KNOWLEDGE BASE.** A collection of information that a human or AI system can draw on to solve problems.

**KNOWLEDGE DOMAIN.** The complete set of knowledge a particular company operates within, as determined by the industry in which it participates.

**KNOWLEDGE ENGINEERING.** The process of revealing the structure of information for best use by humans or artificial intelligence.

**KNOWLEDGE GRAPH.** The relationship between various concepts and data elements in an ontology used to describe the knowledge domain.

**KNOWLEDGE MANAGEMENT SYSTEM.** A system that supports employee productivity by making the best information available to, and easily findable for, employees who need it.

**KNOWLEDGE PORTAL.** A structured knowledge base that supports employee productivity and customer self-service by making the best information available and easily findable in the context of an employee or customer's task.

**KPIS (KEY PERFORMANCE INDICATORS).** Metrics used to measure the level of success of a system.

**LABORATORY INFORMATION MANAGEMENT SYSTEMS (LIMS).** A technology that collects, manages, and analyzes masses of data from laboratory instrumentation.

**LEGACY SYSTEMS.** Older technology systems that an enterprise depends on but that may not be configured to work easily with modern AI algorithms.

**LUMPERS.** People who have a tendency to gather things together into larger groups. *See also* splitters.

**MACHINE LEARNING.** An umbrella terms for AI applications that identify patterns, make predictions, find anomalies, and classify images or data.

Machine learning algorithms solve problems by using increasingly accurate approximations as input to improve the output.

**MARKETING AUTOMATION.** A technology that manages marketing campaigns and customer engagement activities across channels.

**MARKETING QUALIFIED LEAD (MQL).** A lead that meets two of the five BANTS criteria.

**MARKETING TECHNOLOGY STACK.** The collection of interconnected marketing technologies that, together, enable a company to attract and interact with customers.

**MASTER DATA.** A consistent set of identifiers that describe elements of an ontology and how they are applied to enterprise applications.

**MATURITY.** The level of sophistication of a technology system or organization, measured against a rubric.

**MENTAL MODEL.** How a customer or other user thinks about a problem and how to solve it.

**MERCHANDISER.** A person whose job is to make changes in the presentation of products in a retail store or site so as to maximize sales.

**MESSAGING ARCHITECTURE.** A framework that allows for the systematic customization and personalization of marketing messages.

**METADATA.** Data about data or descriptions of data; “is-ness” and “about-ness” of content and information. For example, a video is data, while information about the length, author, format, and date of creation of the video would be called metadata. Also called features, attributes, facets, properties, fields, and columns.

**METRICS-DRIVEN GOVERNANCE.** A framework that allows executives, managers,

and front-line employees to make decisions and execute them based on careful analysis of data and KPIs about the enterprise.

**MODEL-BASED DESIGN.** The process of conceiving, creating, and testing products through virtual simulations.

**MULTIPHYSICS.** A term for simulations that involve multiple models or simulation systems for any given product.

**NATURAL-LANGUAGE PROCESSING (NLP).** Machine algorithms that parse text or speech to generate meaning.

**NAVIGATIONAL HIERARCHY.** A product taxonomy used to aid site navigation. *See also* display taxonomy.

**NET PROMOTER SCORE (NPS).** A simple likelihood-to-recommend metric that organizations use to measure customer satisfaction.

**NEURAL NETWORK.** A computing structure that generates optimal outputs by identifying patterns in broad and varied input data. Neural networks loosely mimic biological processes with artificial “neurons” that compute an approximate output based on changing the weights associated with each element in a network. The output of each “neuron” is used as an input to generate the next iteration of adjustments.

**OMNICHANNEL MARKETING/MESSAGING.** A marketing practice that coordinates messaging to a given customer across channels and devices.

**ONBOARDING.** The process of adding new products to a product information system.

**ONTOLOGY.** The master knowledge scaffolding of an organization: a complete, consistent data model of the customers, products, processes, and relationships that make a business work.

**ORCHESTRATION.** The coordination and synchronization of marketing activities that happens in a mature digital marketing environment.

**PERSONA.** An iconic or archetypal representation of a specific type of customer.

**PERSONALIZATION.** Customization of site content or marketing messages for each individual customer.

**PERSONALIZATION ENGINE.** Software components that determine which data and content should be delivered in each channel to each customer on each device.

**PIM.** *See* product information management system.

**PLANOGRAM.** A diagram showing how products are organized in a retail store.

**POLYHIERARCHY.** A classification system in which some items are located in more than one place.

**PREDICTIVE ANALYTICS.** A process that anticipates future states by analyzing current information and past patterns of behavior.

**PREDICTIVE OFFERS.** An ecommerce technique that tries to anticipate the offer most likely to entice a given customer, based on past actions of similar customers.

**PRIMARY (CUSTOMER) RESEARCH.** The practice of talking to, interviewing, or observing customers in their actual environment.

**PRODUCT INFORMATION MANAGEMENT (PIM) SYSTEM.** A database that holds information about products.

**PRODUCT LIFECYCLE INTELLIGENCE (PLI).** The use of advanced analytics to optimize innovation across all stages of the product lifecycle. The successor to product lifecycle management (PLM).

**PROMISEWARE.** AI tools that promise functionality but don't actually deliver it.

**PROPENSITY-TO-BUY MODEL.** A scoring algorithm that models the likelihood that a prospect will make an actual purchase.

**RACI CHART.** A chart that tracks which entities and people are responsible, accountable, consulted, or informed for any given decision or responsibility.

**REFERENCE DATA.** Standard definitions and values used to ensure consistency of information across systems. The ontology contains reference data for enterprise systems, applications, and processes.

**RIGHTS MANAGEMENT.** A process that ensures that use of content respects the rights of the original creator or licensor of the content.

**ROBOTIC PROCESS AUTOMATION (RPA).** Automation technology that duplicates the activities of a human in moving data among systems.

**ROT.** An acronym for redundant, out-of-date, and trivial content.

**SALES QUALIFIED LEAD (SQL).** A prospect that sales has qualified for three of the five BANTS criteria.

**SCENARIOS.** In the context of an ontology, a description of how a problem is solved as a series of steps. In the context of site design, a description of steps that a particular type of customer might take to solve a problem.

**SCHEMA.** A structure for data, content, or knowledge in a database.

**SEARCH ENGINE OPTIMIZATION (SEO).** The activity of modifying website content to enable pages to appear higher in search engine results.

**SEED DATA.** Initial contexts and examples used to train intelligent virtual assistants.

**SEMANTIC LAYER.** Ontology elements that translate inconsistent information structures into a common language.

**SEMANTIC SEARCH.** A mechanism for integrating diverse structured and unstructured information sources to provide appropriate data to aid in a task leveraging the semantic layer.

**SHOPPING BASKET ANALYSIS.** An ecommerce technique that makes suggestions based on past purchases of customers who have bought similar items.

**SMART OBJECTS.** Sensor-enabled physical goods that can communicate data about themselves.

**SMART SPACES.** Instrumented buildings that track activity within their spaces.

**SOFTWARE AS A SERVICE (SAAS).** A model in which companies subscribe to software services for various technology needs.

**SPLITTERS.** People who have a tendency to separate things into smaller and smaller subcategories. *See also* lumpers.

**STRUCTURED KNOWLEDGE.** Knowledge that is explicitly created for reuse and retrieval from a knowledge base.

**TACIT KNOWLEDGE.** Knowledge that is generated through experience but not documented.

**TAGGING.** The application of metadata attributes to enable the location of data and content relevant to a particular problem or situation.

**TAXONOMY.** A clearly defined hierarchical structure for categorizing information of a given type using parent–child and whole–part relationships. *See also* display taxonomy.

**TECHNICAL DEBT.** Friction and challenges created by technology systems that

were not well documented, kept up-to-date, or integrated with other systems.

**TEXT ANALYTICS.** A set of algorithms that read text and identify patterns or elements within the text, for tagging, sentiment analysis, categorization, and other purposes.

**TEXT MINING.** The process of deriving insights from unstructured collections of text.

**TRAINING DATA.** A subset of a collection of data that an AI algorithm uses to identify patterns.

**UNSTRUCTURED KNOWLEDGE.** Knowledge generated through ongoing collaboration activity that is not designed for reuse.

**UTTERANCE.** A bit of sound or text that a user generates, which is then analyzed by an AI system to extract meaning and intent.

**VIRTUAL AGENT.** A domain-specific, task-oriented, and contextualized algorithm that can respond to user queries. *See also* virtual assistant.

**VIRTUAL ASSISTANT.** An algorithm that uses AI to generate information to make a human's job easier. Includes virtual sales assistants that help customers with ecommerce and intelligent virtual assistants that answer support questions.

**VIRTUAL REALITY.** A technology that displays a digital world that people can move through and interact with, typically on a visual headset.

**VOICE-OF-THE-CUSTOMER.** Term to describe programs that gather customer feedback to provide insights into a business.

## ACKNOWLEDGMENTS

**W**riting *The AI-Powered Enterprise* has been quite a journey over the past two-plus years. I want to thank my wife, Lisa, for putting up with my weekday (and even weekend) evening working sessions and supporting me throughout the process.

This was a team effort, starting with Josh Bernoff who helped me at the start and has continued with extended editorial effort throughout the development process, and Maggie Langrick, the publisher and founder of LifeTree Media, who helped to rework my stream-of-consciousness writing into a structured, logical flow. Kudos to both of you for your hard work and patience. This book would not exist without you. Thanks also to the whole LifeTree Media team, including Sarah Brohman and Tara Tovell.

I am grateful for the help and support of my colleagues at Earley Information Science. In the past few years, we have done a great deal of prototyping, proof of concept development, and technology investigation. My current and former EIS colleagues who contributed to those efforts include Jeannine Bartlett, Dennis Connolly, Eli Cooley, John Dolce, Dino Eliopoulos, Donna Fritzsche, Jeanna Giordano, Tobias Goebel, Gustavo Gonzalez, Prakash Govind, Bryan Kohl, Mac McBurney, Carla Pealer, Vivek Shivaprabhu, Ash Subramanian, and Henry Truong (who also provided the examples from TeleTech). The other members of the EIS team whose hard work make everything we do possible include Mike Anthony, Rachel Benson, Ian Galloway, Jeremy Grubman, Jason Hein (to whom I also

owe the fuzzy socks metadata example), Chantal Schweizer, Dave Skrobela, Melissa Wilkins, and Laura Wright. Special thanks to Sharon Foley.

This book would not have been possible without the fascinating stories of the people doing great AI work and the people who introduced me to them. They were kind to allow me to include them here. These include the folks at PR 20/20, including Mike Kaput, who runs their Marketing AI Institute; Paul Roetzer, founder of the agency and the terrific Marketing AI conference; and Sandie Young, who set up multiple meetings and who introduced me to rasa.io, MarketMuse, and Pandata. I'd also like to thank Cal Al-Dhubaib, Jane Alexander, Mike Barton, the late Darcy Belanger, Jeff Coyle, Nadine Harder, Brett Knight, Jared Loftus, Stephanie McCay, Ryan Miller, Niall Murphy, Melanie Nuce, Molly Soat, Wes Sprinkle, Robert Tercek, and Erik Wolf.

My research and analysis are better because of the smart people I've spoken with over the years. These include Scott Abel, Prith Banerjee, Mathieu Bernard, Ursula Cottone, Brice Dunwoodie, Siobhan Fagan, Jim Gilligan, Jeffrey Hayzlett, Rick Hutton, Carla Johnson, Pete Johnson, Marina Kalika, Gene Kolker, Rebekah Kowalski, Dan Miller, Piero Molino, Mark Nance, Steve Orrin, Steven Keith Platt, Devashish Saxena Kinjal Shah, Laks Srinivasan (who really understood the need for IA in machine learning AI applications), David Talby, Brandon Thomas, Gwen Thomas, Derek Top, Steve Walker, and Rich Wendell. Their insights and stories about analytics, machine learning, chatbots, and AI have inspired me.

A special thank you to Joyce Gavenda, one of my first clients over 25 years ago, whose help over the years has been critical to the company.

I also want to express appreciation for the long-time colleagues who worked on some of our best projects: Jeff Carr, architect for PCL and Applied Materials, who co-developed the IA and ontology development methodologies discussed in *The AI-Powered Enterprise*; Seth Maislin, who contributed to the future scenario that opens the book; Amber Swope and Paul Wlodarczyk, who were instrumental in the success of the Allstate and Applied Materials projects; Branka Kosovac, who worked on many marquis customers; and Stephanie Lemieux, who made contributions to taxonomy methodologies and programs over the course of her tenure with EIS.

Thanks to the hundreds of EIS customers I have had the privilege to work with over the past 30 years and the many people who have passed through the virtual and physical doors of EIS, who are too numerous to mention but have contributed to the lessons, the learning, and the knowledge encompassed in this book.

Looking back, thanks to Gary Kahn, who mentored me early in my career on knowledge management projects with Lotus and then joined EIS to help us become successful.

Finally, thanks to Tom Davenport, not only for writing the foreword and providing candid feedback on the manuscript, but also for participating in EIS webinars and knowledge salons, being part of my research, and making various introductions to folks in the AI space—as well as endangering himself in Falmouth Harbor while experiencing the Thundercat.

# NOTES

## Chapter 1: The Promise and the Challenge of AI

1. This meaningless statement is from “AI Will Bring About the Biggest Transformation in Human History,” Futurism.com, March 9, 2017, <https://futurism.com/ai-will-bring-about-the-biggest-transformation-in-human-history>.
2. For more detail, see David Rotman, “AI is reinventing the way we invent,” *Technology Review*, February 15, 2019, [www.technologyreview.com/s/612898/ai-is-reinventing-the-way-we-invent](http://www.technologyreview.com/s/612898/ai-is-reinventing-the-way-we-invent).
3. I wrote about this in an article: Seth Earley, “There Is No AI Without IA,” *IT Professional* 18, no. 3 (2016): 58–64, <https://ieeexplore.ieee.org/document/7478581>.
4. Here’s a review of Gareth Morgan’s metaphor: Iman Tohidian and Hamid Rahimian, “Bringing Morgan’s metaphors in organization contexts: An essay review,” *Cogent Business & Management* 6, no. 1 (2019). See <https://www.tandfonline.com/doi/full/10.1080/23311975.2019.1587808>.
5. While some progress has been made in having AI fix the data that’s fed into it, the use cases are fairly narrow and significant preparation is needed to make them work. This includes adding “reference data” and numerous examples of what the output should look like.

## Chapter 2: Building the Ontology

1. This description is from the company's home page at [appliedmaterials.com](http://appliedmaterials.com).
2. This example is inspired by the research article from IBM about using multiple specialist bots to solve complex problems: Sethuramalingam Subramaniam and Garbi B. Dasgupta, "COBOTS—A Cognitive Multi-Bot Conversational Framework for Technical Support," in *Proceedings of the AAMAS 2018 Conference* (Stockholm, Sweden, July 10–15, 2018), 597–604 <http://ifaamas.org/Proceedings/aamas2018/pdfs/p597.pdf>.
3. Quote from P.V. Kannan, *The Age of Intent: Using Artificial Intelligence to Deliver a Superior Customer Experience* (Herndon, VA: Amplify, 2019), 26.
4. The game is played by finding the shortest list of movies that link an arbitrarily selected actor and Kevin Bacon. See [https://en.wikipedia.org/wiki/Six\\_Degrees\\_of\\_Kevin\\_Bacon](https://en.wikipedia.org/wiki/Six_Degrees_of_Kevin_Bacon).
5. Albert Einstein is quoted as saying, "If I had an hour to save the world, I would spend 59 minutes identifying the problem and 1 minute solving it," in Dwayne Spradlin, "Are You Solving the Right Problem?" *Harvard Business Review*, September 2012, <https://hbr.org/2012/09/are-you-solving-the-right-problem>.

## Chapter 3: Customer Experience: The Front Line of the Battle

1. Based on Jill Avery, Susan Fournier, and John Wittenbraker, "Unlock the Mysteries of Your Customer Relationships," *Harvard Business Review*, July–August 2014, <https://hbr.org/2014/07/unlock-the-mysteries-of-your-customer-relationships>.

## Chapter 5: Making Ecommerce Smarter

1. Jason Hein used this example in a conference keynote address.
2. According to Salesforce, personalized product recommendations boost revenue. Statistics from Heike Young, "Personalized Product Recommendations Drive 7% of Visits but 26% of Revenue," *Salesforce Blog*, November 2, 2017, <https://www.salesforce.com/blog/2017/11/personalized-product-recommendations-drive-just-7-visits-26-revenue.html>.
3. Quoted in Jared Spool, "When It Comes to Personas, The Real Value Is In The Scenarios," UIE, September 11, 2018, <https://articles.uae.com/when-it-comes-to-personas-the-real-value-is-in-the-scenarios>.

### Chapter 6: Sharpening the Sales Process

1. A short list of such tools appears in Mike Kaput, “4 Scary Smart AI Tools That Will boost Sales Productivity,” Marketing Artificial Intelligence Institute, October 8, 2018, <https://www.marketingaiinstitute.com/blog/ai-for-sales>.
2. The EverString claim is documented in “The World’s Best Data,” a marketing piece posted at [http://xd07g309cu21kppg63qolt4t-wpengine.netdna-ssl.com/wp-content/uploads/2018/01/EverString\\_WorldsBestData.pdf](http://xd07g309cu21kppg63qolt4t-wpengine.netdna-ssl.com/wp-content/uploads/2018/01/EverString_WorldsBestData.pdf).

### Chapter 7: Customer Service: Delivering Higher Quality at a Lower Cost

1. Data from the “Customer Service Representatives” page on Deloitte’s Data USA website: <https://datausa.io/profile/soc/customer-service-representatives>.
2. See IA Staff, “70% of Buying Experiences are Based on How the Customer Feels They are Being Treated,” Industry Analysts, Inc., December 5, 2017, <https://www.walkerinfo.com/knowledge-center/featured-research-reports/customers-2020-a-progress-report>; Glance, *Counting the Customer: The Complete Guide to Dynamite Customer Care*, [http://ww2.glance.net/wp-content/uploads/2015/07/Counting-the-customer\\_-\\_Glance\\_eBook-4.pdf](http://ww2.glance.net/wp-content/uploads/2015/07/Counting-the-customer_-_Glance_eBook-4.pdf).
3. Adam Cheyer remarks are from BigSpeak Speakers Bureau, “Adam Cheyer—Siri, Back to the Future,” YouTube video, 55:37, April 27, 2016, See <https://www.youtube.com/watch?v=UBHgj9TuHXM>.
4. Ontologies such as this can use a specific structural system called DITA—Darwin Information Typing Architecture. DITA is a component architecture, developed originally by technical documentation professionals to make information more consumable by computers and more accessible for people to use in various contexts. Once the knowledge is in this format, it can fuel an AI system, which can pick out the exact information a customer is looking for.
5. For more on the ongoing challenges with training Watson, see Jason Bloomberg, “Is IBM Watson A ‘Joke’?” *Forbes*, July 2, 2017, <https://www.forbes.com/sites/jasonbloomberg/2017/07/02/is-ibm-watson-a-joke/#4dfb0577da20>.
6. This is from P.V. Kannan, *The Age of Intent: Using Artificial Intelligence to Deliver a Superior Customer Experience* (Herndon, VA: Amplify, 2019), 143.
7. More detailed analysis in Susan Hash, “Chatbots in the Call Center,” *Contact Center Pipeline*, March 2019, <https://www.contactcenterpipeline.com/Article/chatbots-in-the-contact-center>.

## Chapter 8: Accelerating Employee Productivity

1. TeleTech changed its name to TTEC in 2018.

## Chapter 9: Physical Meets Digital: Manufacturing, Supply Chain, and Logistics

1. Thanks to Cal Al-Dhubaib, Managing Partner at Pandata, for making us aware of the work of the Cleveland Museum of Art.
2. A description of digital twins appears on the “Digital Twin” page on Siemens’ website: <https://www.plm.automation.siemens.com/global/en/our-story/glossary/digital-twin/24465>.
3. The study appeared in the IEEE transactions journal: S.H. Huang and Hong-Chao Zhang, “Artificial neural networks in manufacturing: concepts, applications, perspectives,” *IEEE Transactions on Components, Packaging, and Manufacturing Technology: Part A* 17, no. 2 (1994): 212–228, <https://ieeexplore.ieee.org/document/296402>.
4. More detail in Robert Gebelhoff, “Sequencing the Genome Creates So Much Data We Don’t Know What To Do with It,” *Washington Post*, July 7, 2015, [https://www.washingtonpost.com/news/speaking-of-science/wp/2015/07/07/sequencing-the-genome-creates-so-much-data-we-dont-know-what-to-do-with-it/?utm\\_term=.6585c7ec7969](https://www.washingtonpost.com/news/speaking-of-science/wp/2015/07/07/sequencing-the-genome-creates-so-much-data-we-dont-know-what-to-do-with-it/?utm_term=.6585c7ec7969).
5. For more on simulations, see Carlos Gonzalez, “The Impact of Simulation and the Future of Manufacturing,” *Machine Design*, February 1, 2017, <https://www.machinedesign.com/fea-and-simulation/impact-simulation-and-future-manufacturing>.
6. Quote from Insights Team, “How AI Builds A Better Manufacturing Process,” *Forbes Insights*, July 17, 2018, <https://www.forbes.com/sites/insights-intelai/2018/07/17/how-ai-builds-a-better-manufacturing-process/#6b19aff41e84>.
7. The video describing this is Jordan Reynolds and Ryan Whittle, “KNOW/The Future of Product Data Analytics,” *Viewpoints on Innovation*, January 10, 2019, <http://viewpoints.io/entry/know-the-future-of-product-data-analytics>.
8. Quote from Luigi Alberto Ciro De Filippis, Livia Maria Serio, Francesco Facchini, and Giovanni Mummolo, “ANN Modelling to Optimize Manufacturing Process,” *IntechOpen*, December 20, 2017, <https://www.intechopen.com/books/advanced-applications-for-artificial-neural-networks/ann-modelling-to-optimize-manufacturing-process>.

9. For more insights, see Zohar Yami, Golan Meltser, and Rotem Greener, “Supply Chain Complexity—Dealing with a Dynamic System,” *Tefen Tribune*, Spring 2010, [http://tefen.com/uploads/insights/1456214433\\_ZyuFmdCdXJ.pdf](http://tefen.com/uploads/insights/1456214433_ZyuFmdCdXJ.pdf).
10. Cited from Howard Slusken, “Four million parts, 30 countries: How an Airbus A380 comes together,” CNN, December 2018, <https://www.cnn.com/travel/article/airbus-a380-parts-together/index.html>.
11. Cited from Jonathan Webb, “How Many Suppliers Do Businesses Have? How Many Should They Have?” *Forbes*, February 28, 2018, <https://www.forbes.com/sites/jwebb/2018/02/28/how-many-suppliers-do-businesses-have-how-many-should-they-have/#1ad2fbc9bb72>.

### Chapter 10: AI-Powered Strategy and Governance

1. More on this in Jen Booton, “Under Armour Launches Ecosystem of Connected Fitness Products,” *SportTechie*, January 16, 2019, <https://www.sporttechie.com/under-armour-launches-ecosystem-of-connected-fitness-products>.
2. More on this in Kate Grosch, “John Deere—Bringing AI to Agriculture,” *HBS Digital Initiative*, November 7, 2018, <https://digital.hbs.edu/platform-rctom/submission/john-deere-bringing-ai-to-agriculture>.

### Chapter 11: Leading into the Future

1. Data from Evan Andrews, “10 Things You May Not Know About the Apollo Program,” History.com, September 1, 2018, <https://www.history.com/news/10-things-you-may-not-know-about-the-apollo-program>.
2. For more detail, see Burton Dicht, “The Greatest Engineering Adventure Ever Taken,” ASME, December 28, 2010, <https://www.asme.org/engineering-topics/articles/history-of-mechanical-engineering/the-greatest-engineering-adventure-ever-taken>.
3. For more detail, see Rhuaridh Marr, “To the Moon and Back on 4KB of Memory,” *Metro Weekly*, July 24, 2014, <https://www.metroweekly.com/2014/07/to-the-moon-and-back-on-4kb-of-memory>.
4. For more detail, see Sasjia Otto, “Apollo 11 moon landing: top 15 Nasa inventions,” *The Telegraph*, July 22, 2009, <https://www.telegraph.co.uk/news/science/space/5893387/Apollo-11-moon-landing-top-15-Nasa-inventions.html>.
5. This refers to Robert Tercek’s book *Vaporized: Solid Strategies for Success in a*

*Dematerialized World* (Vancouver: LifeTree Media, 2015).

6. More detailed analysis available in Brian Solis, “The State of Digital Transformation,” Altimeter, 2019, <https://insights.prophet.com/the-state-of-digital-transformation-2018-2019>.
7. Patti Anklam, an expert and consultant in ONA, wrote an informative book on the topic: *Net Work: A Practical Guide to Sustaining Networks at Work and in the World* (Oxford, UK: Butterworth-Heinemann, 2007).

# INDEX

- [24]7.ai, 157
- 360-degree view of customer, 83
- ABIE, 145–47, 149, 150, 151, 154, 183
- about-ness, 33, 42–43, 174, 176, 273
- Age of Intent, The* (book), 34, 158
- agility, 13–14, 14–16, 241, 268
  - problems that prevent, 14–16
- AI (artificial intelligence), 9–12, 1–6, 273
  - AI roadmap, 244–45
  - AI-powered customer experience, 58–74
  - analyzing manufacturing with, 202–3
  - assessing projects realistically, 261–63
  - budgets for, 258–59
  - cognitive, 24–25
  - evaluating salespeople with, 137–42
  - for knowledge management, 174–77
  - fragmented initiatives, 260–61
  - helps drive strategy, 215–18
  - helps salespeople, 123–24
  - hiring for, 259
  - in biology, 197–98
  - in content marketing, 81–83
  - in hiring, 185–87
  - in mining, 213
  - in product lifecycle, 198–202
  - in SEO, 75–78
  - in site customization, 104
  - ontology powers, 33–35
  - preparing for, 238–39
  - priorities for, 267
  - risks of, 267
  - strategy, 244–245
  - surfaces high-value content, 178–79
  - tips for leaders, 266–72
  - tools often fail, 70
  - vendors overpromise, 259
  - why projects fail, 258–60
  - will determine success, 238–39
- AI Advantage, The* (book), 11, 14
- Airbnb, 14
- Airbus, 206
- Alexa, 6
- Alexander, Jane, 189–91, 204
- Allstate Business Insurance, 144–47, 150, 155
- Altimeter Group, 264
- Amazon Lex, 271
- American Marketing Association, 81–83, 84
- apocalypse, robot, 272
- Apollo program, 238
- Applied Materials, 27–30, 35–36
  - ontology at, 35–36
- artificial intelligence. *See* AI
- ArtLens Wall, 189
- as-built records, 200
- attribute model, 273
- audiences, 41, 97, 273. *See also* personas
  - attributes of, 100–101
  - customer taxonomy, 99
- augmented reality, 191, 195, 274
- Autodesk, 200
- Avery, Jill, 64

- B2B, 108–9, 114–15
- Bank of America, 3
- BANTS criteria, 125, 274
- Barton, Mike, 144, 145, 147
- beacons, 212
- big data, 153, 274
- biology, AI in, 197–98
- body language. *See* digital body language
- Borders, 1
- born digital, 14, 239
- bots, 31, 34, 274
  - helper bots, 157
- CAD/CAM, 194, 274
- call center, 138, 144–47
  - at Allstate Business Insurance, 144–47
- Cas9, 198
- categorization, 41, 43
- CDP. *See* customer data platforms
- change management, 271
- chatbots, 3–6, 124–25, 149, 158–60, 237, 274
  - in customer service, 158–59
  - in sales, 124–25
- Cheyser, Adam, 149
- Chief Data Officer (CDO), 243
- Cleveland Museum of Art, 189–91, 204, 211
- clock speeds, variation within enterprises, 21–23
- cobots (collaborative robots), 195, 274
- cognitive AI, 24–25
- cognitive load, 55, 56, 92, 182, 275
- collaboration site, 170, 171
- communication
  - holistic, 14–15
  - in sales, 127
- Configure-Price-Quote (CPQ) system, 132–34, 135–37, 275
  - checklist for deploying, 135–37
- content architecture, 212
- content management system (CMS), 67, 102, 110, 111, 126, 275
- content marketing, 81–83, 275
- content model, 150, 212, 275
- content quality, 177–78
- context, 150, 275
  - in knowledge management, 173–74, 182
  - in ontology, 183
- contracts, automating, 134, 137
- conversational commerce. *See* conversational marketing
- conversational marketing, 124, 275
- Conversica, 128
- CRISPR, 198
- crowdsourcing, 158, 275
- culture, 242
- current state maturity. *See* maturity
- customer attribute models, 59, 105, 275
- customer data, 97–99
  - audience attributes, 100–101, 102
  - explicit and implicit, 100, 101
- customer data platforms (CDP), 101–2, 130, 276
- customer engagement strategy, 62–64
- customer experience (CX), 11–13, 50–75, 276
  - customer engagement strategy, 62–64
  - customer journey, 57–58
  - customer lifecycle, 60–62
  - defined, 52
  - digital marketing in, 78–79
  - high-fidelity journey map, 58–74
  - journey map, 57, 58
  - knowledge management supports, 181
- customer feedback, 160–62
- customer journey, 12, 56, 57–58, 132, 276, 278. *See also* journey map
  - knowledge supports, 179–81
- customer lifecycle, 60–62, 241, 276

- customer model, 62, 276
- customer relationship management
  - system (CRM), 67, 102, 106, 126, 127, 139–42, 242, 276
  - compliance issues with, 127
  - ontology connects with, 139–42
- customer service, 144–62
  - at unicorn startup, 266
  - costs of, 147
  - FAQs in, 148
- customization, 102–6. *See also*
  - personalization
  - AI helps determine, 104
  - at scale, 104–6
- data. *See also* data quality, training data
  - audience attributes, 100–101
  - customer data, 97–99
  - data-driven mindset, 242–44
  - incompatible languages for, 15
  - junk data, 16
  - organizing principles for, 42–43
  - powers ecommerce, 93–102
  - product data, 94–97
  - role in AI-powered enterprise, 17, 18–19
  - tagging, 16–17
- data dictionaries, 37, 276
- data hygiene, 196, 276
- data models, 58, 276
- data quality, 37, 84, 112, 116, 178, 241, 277. *See also* junk data
  - who is responsible, 241
- Davenport, Tom, 9–12, 14
- design win, 121
- digital asset management (DAM) system, 84, 190, 277
- digital body language, 79, 84, 105
- digital marketing, 75–86
  - agile, 79–81
  - central to customer experience, 78–79
  - content marketing, 81–83
  - is fragmented, 79
  - marketing automation, 84, 281
  - marketing technology stack, 65, 83
  - messaging architecture, 104
  - new roles for, 83–85
  - omnichannel messaging, 80
  - ontology in, 85
  - orchestrating messages, 80–81
- digital transformation, 124, 239, 240, 241, 263–66, 277
  - maturity assessment for, 246–47
- digital twin, 194–95, 199, 277
  - in product testing, 194
- display taxonomy, 94, 277
- disruption, 267
- DITA (Darwin Information Typing Architecture), 229, 276
- domain model, 161, 277
- dynamic navigation, 91–92
- ecommerce, 86–120
  - assessing maturity, 111–19, 114–15
  - B2B, 108–9
  - conversational. *See* conversational marketing
  - data powers, 93–102
  - dynamic navigation, 91–92
  - ontology in, 106–8
  - predictive offers, 92
  - sample ontology, 108
  - shopping basket analysis, 93
  - site customization, 102–6
  - terminology used in, 99–100
  - website navigation, 89–93
- employee journey, 178, 277
- engagement. *See* customer engagement
  - strategy
- enterprises, organic view of, 9–11
- entity extraction, 91, 277
- entropy, 16
- equivalence terms, 107
- ERP (enterprise resource planning), 21, 29, 39, 102, 106, 126, 277
- EverString, 129
- EVERYTHING, 207–8, 229
- Expedia, 5

- explicit descriptors, 62, 278
- facets, 90, 278
- features, in AI, 106, 174, 278
- FIFSG, 156–57, 158
- fitness to purpose (applications), 110
- Fournier, Susan, 64
- friction, 13, 15, 28, 37, 53, 64, 179, 278
- Get Smart, 167
- Glengarry Glen Ross, 124
- Godwin, Kim, 98
- Google, 14, 89
  - Dialogflow, 271
- governance, 23, 221–34
  - bodies that manage, 226–29
  - establishing a framework for, 223
  - establishing a framrework for, 225–34
  - gaps in, 224
  - managing resources, 229, 232
  - managing technology, 233–34
  - metrics-driven, 221–34, 269
  - negative connotations of, 223
  - reporting tools for, 229
  - training, 234
- Grainger, 4
- graph data, 129
- groupware, 165, 278
- GS1, 209, 220, 229
- helper bots, 157, 278
- heuristics, 245, 278
- high-fidelity journey maps, 12, 58–74, 279
  - defined, 58
  - must be validated, 59
  - steps in creating, 60–74
- Home Depot, 1
- Honeywell, 220
- HR, AI in hiring, 185–87
- HubSpot, 131
- IA. *See* information architecture
- IBM, 6, 155, 186
- implementation roadmap, 71–74
- implicit descriptors, 62, 279
- inference engine, 34, 152, 279
- information architecture (IA), 9, 8, 35, 36, 150, 244, 266, 279
- information ecosystem, 204
- information flows, 164–65, 264
  - customer data, 102
  - in manufacturing, 193
  - in organic view of enterprise, 10, 11
  - in supply chains, 204, 205
  - knowledge management improves, 164–65
  - unclogging, 16–17
- information hygiene, 168
- information metabolism, 31, 37, 279
- infrastructure programs, 12, 279
- Ingersoll-Rand, 132
- injection molding, 202-3
- innovation, 193
- intelligent virtual assistant. *See* virtual assistant
- intent, 34, 279
- internet of things (IoT), 218
- intranet, 163, 166
- is-ness, 33, 42–43, 174, 176, 279
- IT (information technology) department, 241–42
- Jeopardy, 155
- John Deere, 220
- journey map, 57, 58–74, 58, 279. *See also* high-fidelity journey maps
- Juengst, David, 77
- junk data, 16
- Kalypso, 200–201
- Kamikoto, 51–52, 62
- Kannan, P.V., 34, 157
- Knight, Brett, 70
- knowledge architecture, 35, 279
- knowledge base, 34, 117, 145, 146, 148, 153, 157, 166, 175, 176–77, 179–81, 196, 280
  - automation helps fill, 176–77
  - in manufacturing, 196
  - supports the business, 179–81
- knowledge engineering, 153, 154–56, 164, 280

- knowledge graph, 34, 280
- knowledge management, 165, 168–79
  - AI helps with, 174–77
  - content quality, 177–78
  - context in, 173–74, 182
  - explicit vs. tacit knowledge, 175
  - high vs. low-value knowledge, 171–73, 178–79
  - knowledge management system, 168–70, 280
  - maturity model for, 254
  - measuring effectiveness, 183–85
  - ontology provides consistency, 176
  - structured vs. unstructured knowledge, 169–71, 174
  - supports customer journey, 179–81
  - supports the business, 179–81
- knowledge portal, 151, 152, 280
- Kormushoff, Alex, 215
- laboratory information management systems (LIMS), 197, 280
- Leaderscape, 215–17, 217
- leads
  - capturing, 124–25
  - communication cadence and messaging, 127
  - nurturing, 127–28
  - prioritizing, 128–31
- legacy systems, 20, 280
- lifecycle. *See* customer lifecycle
- Lotus Notes, 165
- lumpers, 43, 280
- machine learning, 10, 24, 105, 106, 107, 129, 153, 174, 197, 202, 280
- manufacturing, 192–98
  - AI analysis of, 202–3
  - AI in product design, 200
  - as-built records, 200
  - company using paper catalogues, 121–23
  - data hygiene in, 196
  - digital twins in, 194–95
  - model-based design, 199
  - ontology in, 193
  - product innovation, 193
  - product lifecycle, 198–202
  - product lifecycle intelligence, 200–202
  - quality control, 196–97
  - rework decision, 199–200
  - robotic, 195–96
- marketing. *See also* digital marketing
  - conversational, 124, 275
- marketing automation, 84, 281
- marketing qualified lead (MQL), 125, 281
- marketing technology stack, 65, 83, 281
- MarketMuse, 75–78
- Marketo, 131
- master data, 19, 106, 281
- Mathews, Brian, 200
- maturity, 245–56, 269, 281
  - assessing, 65–66, 111–19, 114–15, 246–47, 247–58
  - benchmarking, 247–58
  - checklist for, 113, 116–19
  - ecommerce, 111–19, 114–15
  - five-stage model for, 245–56
  - for content optimization, 248
  - for customer information, 250
  - for knowledge management, 254
  - for orchestration, 256
  - for product information, 252
  - maturity models, 247–58
- Medicare, 263
- mental models, 30, 36, 62–64, 99
  - customization can break, 92
- messaging architecture, 104, 281
- metadata, 8, 62, 106, 107, 110, 281
- metrics-driven governance, 221–34, 269, 281
- Microsoft
  - HoloLens, 191
  - LUIS, 271
- Miller, Ryan, 156–57
- mining, 213
- model-based design, 199, 282
- mold stripping, 90, 99, 107

- moonshot, 7, 238, 258, 259, 270
- Morgan, Gareth, 9
- Mowi, 208
- MQL. *See* marketing qualified lead
- MSCDirect.com, 90
- multiphysics, 199, 282
- Murphy, Niall, 208
- MyFitnessPal, 220
- natural-language processing (NLP), 127, 282
- navigational hierarchy, 94, 282
- navigational taxonomy, 96
- NCR Teradata, 155–56
- nesting, 138
- Net Promoter Score, 161, 282
- neural networks, 104, 202, 266, 282
- omnichannel, 80, 117, 282
- onboarding, 93, 111, 116, 282
- ontology, 8–9, 23–24, 27–49, 282
  - accidental, 38, 273
  - and cognitive AI, 24–25
  - applying, 47–48
  - as “Rosetta Stone”, 31, 38, 99
  - as North Star, 37
  - as reference for knowledge architecture, 176
  - as reference librarian, 33, 35, 107
  - as source of truth, 36–37, 107, 177
  - at Allstate Business Insurance, 146, 155
  - at Applied Materials, 35–36
  - attributes in, 107–8
  - audiences’ role in, 41
  - cannot build haphazardly, 38–39
  - context in, 183
  - data- and content-centric approach to, 43–45
  - defined, 8–9
  - design elements in, 160
  - drives search, 92
  - enables standards, 209–10
  - equivalence terms, 107
  - explained, 32–33
  - graph data, 129
  - how to build, 40–47
  - in digital marketing, 85
  - in ecommerce, 106–8
  - in knowledge management, 177
  - in manufacturing, 193
  - in semantic search, 131
  - in supply chains, 207
  - like chart of accounts, 29, 32
  - powers AI, 33–35
  - powers sales AI, 139–42
  - reconciles agility and scalability, 54
  - ROI of, 14
  - sample ecommerce ontology, 108
  - scenarios in, 41
  - top-down approach to building, 40–44
  - use cases in, 41
  - used for maintenance, 212–13
  - validating, 45–47
  - varies by industry, 39
  - why ontologies matter, 30–32
- operationalization, 18, 21–23
- orchestration, 80, 268, 283
- paper catalogues, 121–23
- paradox of choice, 56
- PCL Construction, 163–64, 164, 175, 179
- Perkins, Allen, 3–6
- personalization, 85, 118, 236, 283
- personas, 97–99, 237, 283
- PIM. *See* product information management system
- planogram, 95, 283
- polyhierarchy, 92, 283
- Polytechnic University of Bari, Italy, 202
- predictive analytics, 32, 283
- predictive offers, 92, 283
- preventative maintenance, 211
- primary research, 57, 59, 98, 283
- procurement, 208
- product data, 94–97
  - product relationships, 95–96
  - maturity, 252
  - standardized or differentiated, 96–97
- product design, 200

- product information management (PIM)
  - system, 95, 96, 106–7, 107, 111, 126, 283
- product lifecycle, 198–202
  - product lifecycle intelligence (PLI), 200–202, 283
- promiseware, 70, 284
- propensity-to-buy model, 130, 284
- proposals, automating, 134, 137
- publishing company, 222–34
- quality control, 196–97
- RACI chart, 228–29, 230–31, 284
- Ralph Lauren, 208
- Rasa.io, 81, 82
- research. *See* primary research
- Reynolds, Jordan, 200
- rights management, 32, 284
- robot apocalypse, 272
- robotic manufacturing, 195–96
- robotic process automation (RPA), 243, 284
- Rosetta Stone (ontology), 31, 38, 99
- ROT (bad data), 178, 284
- sales, 121–43
  - automating proposals/contracts, 134, 137
  - chatbots in, 124–25
  - Configure-Price-Quote system (CPQ), 132–34, 135–37
  - lead capture, 124–25
  - lead nurturing, 127–28
  - nesting, 138
  - ontology powers AI, 139–42
  - prioritizing leads, 128–31
  - propensity-to-buy model, 130
  - sales qualified lead, 125
  - salespeople, 123–24, 137–42
  - semantic search helps with, 131–32
- sales qualified lead (SQL), 284
- salesforce.com, 127
- salespeople, 123–24, 137–42
  - AI in simulations, 138
  - compliance issues with, 127
  - evaluating with AI, 137–42
  - how AI helps, 123–24
- scalability, 54
- scenarios, 41, 98–99, 237, 263, 284
  - for ontology, 41
  - in customer experience design, 99
- schemas, 152, 284
- search, 89–91, 151, 152
  - as a conversation, 89
  - faceted, 90
  - how ontology drives, 92
  - semantic search, 131–32, 157, 170
- search engine optimization (SEO), 75–78, 284
- seed data, 153, 284
- semantic search, 131–32, 157, 170, 285
- semiconductor manufacturing, 27–30
- SEO. *See* search engine optimization
- Seventh Sense, 131
- shopping basket analysis, 93, 285
- Siri, 149, 171
- Skype, 166
- Slack, 166
- smart objects, 210–11, 285
  - enable predictive maintenance, 211
- smart spaces, 211–12, 285
- Soat, Molly, 81, 82, 84
- social media, 21
- software as a service, 240, 285
- speed, impediments to, 14–16
- splitters, 43, 285
- SQL. *See* sales qualified lead
- standards, 209–10, 232–33
  - standards committee, 232–33
- strategy, 215–18
- supply chains, 203–9
  - AI can support, 207
  - distribution and data, 206–9
  - information ecosystem for, 204
  - information flows in, 204, 205
  - information supply chain, 205
  - ontology essential to, 207
  - traceability of products, 208
- tacit knowledge, 175, 285

- tagging, 16–17, 174, 285
- taxonomies, 9, 28, 33, 39, 237, 285
  - at Allstate Business Insurance, 146
  - at Applied Materials, 35–36
  - defined, 33
  - display taxonomy, 94
  - for customer feedback, 161
  - of customers, 99
  - of service issues, 266
  - product data, 94–97
  - standardized or differentiated, 96–97
  - uses of, 33
  - vary by industry, 39
- TCO (total cost of ownership), 136
- technical debt, 20, 243, 286
- technology
  - analyzing deployment of, 67–69
  - governance for, 233–34
  - prioritizing, 71–74
  - role in AI-powered enterprise, 18, 20
- technology stack, 260. *See also* marketing technology stack
- TeleTech, 185–87
- Teradata, 155–56
- Tercek, Robert, 247
- terminology, 131
- Tesla, 14, 219
- Texas Tech, 196
- text analytics, 29, 91, 107, 134, 160, 161, 174, 176, 178, 196, 286
- Thermo Fisher Scientific, 197, 221
- Tower Records, 1
- training, 234
- training data, 8, 128, 151, 176, 260, 286
- Truong, Henry, 185–87
- Twilio, 166
- Uber, 14
- Under Armour, 220
- unicorn startup, 266
- USAA, 70
- use cases, 41, 46, 263. *See also* scenarios
  - in ontology, 41
  - must be testable, 46
- user-generated content, 117
- vaporization, 247
- virtual agents, 145–47, 151, 152, 286
  - at Allstate Business Insurance, 145–47
  - justifying, 156–60
  - steps to roll out, 158–160
- virtual assistants, 31, 34, 125–27, 149–54, 151, 152–53, 286
  - for customer service, 148, 149–54
  - for information retrieval, 152–53
  - for sales, 87–88, 125–27
- virtual reality, 195, 212–13, 286
- Vodafone, 259, 270
- voice-of-the-customer, 161, 286
- Wall Street Journal, The*, 1
- Walmart, 137
- Watson, 6, 155
- Wayfair.com, 103
- weak signals, 59, 218
- Whittle, Ryan, 200
- WidgetCo, 121–23
- wiki, 166, 171
- Wittenbraker, John, 64
- XML, 233



## ABOUT THE AUTHOR

**T**hroughout his career, Seth Earley has been passionate about the crucial role that information management would play in a world hurtling toward digital transformation. He provides challenging insights to executives who are tasked with leading their organizations forward in an age in which the digital experience offered to customers determines the winner.

As CEO of Earley Information Science, a consulting firm he founded more than 25 years ago, Seth guides some of the world's most recognized brands on how to leverage their information assets to deliver state-of-the-art customer experiences through integrated enterprise architectures. Seth has a long history of contributing to industry education and research in emerging fields. His current work contributes to a better understanding of topics including cognitive computing, knowledge engineering, data management systems, taxonomy, ontology, and metadata governance strategies.

He coined the phrase “There’s no AI without IA,” which calls out the need for a foundational information architecture for AI projects. His phrase was repeated by IBM CEO Ginni Rometty at the World Economic Forum in Davos, Switzerland, in 2019, when she was interviewed about the challenges of AI.

Seth Earley is a sought-after speaker, writer, and influencer. His writing has appeared in IEEE’s *IT Professional* magazine, where, as former editor, he wrote a regular column on data analytics and information access issues and trends. He has also contributed to the *Harvard Business Review*, *CMSWire*, and *Applied Marketing Analytics*. He co-authored *Practical Knowledge Management* from IBM Press.



## ABOUT EARLEY INFORMATION SCIENCE

Earley Information Science is a professional services firm. For 25 years, we have made it our mission to support business outcomes by organizing data—making it findable, usable, and valuable.

Earley Information Science specializes in structuring and organizing enterprise information with service offerings around four pillars: Product Data Management (PD), Knowledge Engineering (KE), Content Optimization (CO), and Customer Engagement (CE).

Our service offerings include:

- **AI strategy and roadmap** (all pillars). Learn how to best leverage AI for business value.
- **Current state maturity** (all pillars). Understand your current state to develop a realistic roadmap and action plan.
- **Conversational commerce readiness** (PD). Prepare for conversational commerce; don't get left behind.
- **Chatbot proof of capability** (PD, KE). As you evaluate chatbots, separate hype from reality.
- **Product data optimization** (PD). Optimize online catalogs for ecommerce and build the correct foundation for advanced capabilities.
- **Readiness for personalization** (all pillars). Review current state of readiness for marketing and ecommerce personalization.
- **Knowledge architecture for AI** (KE). Separate hype from reality, solve problems today while preparing for the future.

- **Metrics-driven governance** (PD, CO, CE). Measure your return on investments, build a metrics-driven playbook, show linkage to business value.
- **Configure/price/quote** (all pillars). Improve efficiencies, reduce overhead, speed responsiveness, and remain competitive.
- **Technology selection** (all pillars). Control the procurement process, don't be manipulated by the vendor sales process, and make decisions that will advance your career.
- **Taxonomy and ontology design** (all pillars). Realize the value of next generation information architecture to support AI and legacy technologies (since there's no AI without IA).

Looking for help? Contact us at +1-781-812-5551 or [info@earley.com](mailto:info@earley.com).