

O'REILLY®

AI Value Creators

Beyond the Generative AI User Mindset



Rob Thomas,
Paul Zikopoulos
& Kate Soule

"A handbook for the AI Renaissance to help entrepreneurs and innovators drive AI value creation at the next level."

will.i.am, founder and CEO, FYI.AI

"Rob Thomas brings insight, common sense, and his long experience at IBM to bear on the greatest technological transformations of our lifetime. On the subject of AI, there are few people whose perspective I would value more."

Malcolm Gladwell, host of the *Revisionist History* podcast

"This handbook provides actionable insights to help you drive innovation and navigate the next wave of AI advancements, positioning your business for long-term success."

Jessica Sibley, CEO, TIME

AI Value Creators

We've arrived in a new era—GenAI and agentic AI are reshaping industries and decision-making processes across the board. As a result, understanding their potential and pitfalls has become crucial. But in order to stay ahead of the curve, you'll need to develop fresh perspectives on leveraging AI beyond mere technical know-how. Geared to business leaders and tech professionals alike, this book demystifies the strategic integration of AI into business practices, ensuring you're equipped not just to participate but to lead in this new landscape.

This insightful guide by industry leaders Rob Thomas, Paul Zikopoulos, and Kate Soule, with contributions from Rebecca Reyes, David Cox, and Linda Snow, goes beyond the basics, offering real-life success stories and learned lessons to provide a blueprint for meaningful AI engagement. Whether you're a novice or an expert, you'll come away with an enhanced understanding of all the things a modern AI strategy can do for your business.

- Recognize the transformative potential of AI in business and how to harness it
- Navigate the ethical and operational challenges posed by AI with confidence
- Understand the interplay between AI technology and business strategy through detailed use cases
- Implement actionable strategies to integrate AI into your organizational culture
- Step confidently into the role of an AI Value Creator, equipped to lead

Rob Thomas is SVP and CCO at IBM and leads its entire software business, including product management, design, and development.

Paul Zikopoulos is an IBM VP focused on skills and AI. He's also an award-winning writer and speaker who has discussed AI and big data with *60 Minutes* and NATO.

Kate Soule is an IBM research director who leads technical product management for Granite, IBM's family of large language models.

AI / DATA

US \$79.99 CAN \$99.99

ISBN: 978-1-098-16834-6



9

O'REILLY®

Praise for *AI Value Creators*

A handbook for the AI Renaissance to help entrepreneurs and innovators drive
AI value creation at the next level.

—*will.i.am, founder and CEO, FYI.AI*

Rob Thomas brings insight, common sense, and his long experience at IBM to bear on
the greatest technological transformations of our lifetime. On the subject of AI, there are
few people whose perspective I would value more.

—*Malcolm Gladwell, host of the Revisionist History podcast*

With AI reshaping industries, this handbook provides actionable insights that can help
you drive innovation and navigate the next wave of AI advancements, positioning your
business for long-term success.

—*Jessica Sibley, CEO, TIME*

AI Value Creators

Beyond the Generative AI User Mindset

Rob Thomas, Paul Zikopoulos, and Kate Soule

O'REILLY®

AI Value Creators

by Rob Thomas, Paul Zikopoulos, and Kate Soule

Copyright © 2025 O'Reilly Media, Inc. All rights reserved.

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (<http://oreilly.com>). For more information, contact our corporate/institutional sales department: 800-998-9938 or corporate@oreilly.com.

Acquisitions Editor: David Michelson

Development Editor: Gary O'Brien

Production Editor: Kristen Brown

Copyeditors: Doug McNair and nSight, Inc.

Proofreader: Sonia Saruba

Indexer: Potomac Indexing, LLC

Interior Designer: David Futato

Cover Designer: Karen Montgomery

Illustrator: Kate Dullea

April 2025: First Edition

Revision History for the First Edition

2025-04-01: First Release

See <http://oreilly.com/catalog/errata.csp?isbn=9781098168346> for release details.

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. *AI Value Creators*, the cover image, and related trade dress are trademarks of O'Reilly Media, Inc.

The views expressed in this work are those of the authors and do not represent the publisher's views. While the publisher and the authors have used good faith efforts to ensure that the information and instructions contained in this work are accurate, the publisher and the authors disclaim all responsibility for errors or omissions, including without limitation responsibility for damages resulting from the use of or reliance on this work. Use of the information and instructions contained in this work is at your own risk. If any code samples or other technology this work contains or describes is subject to open source licenses or the intellectual property rights of others, it is your responsibility to ensure that your use thereof complies with such licenses and/or rights.

This work is part of a collaboration between O'Reilly and IBM. See our [statement of editorial independence](#).

978-1-098-16834-6

[LSI]

Table of Contents

Preface.....	xi
1. +AI to AI+: Generative AI and the “Netscape Moment”.....	1
What Is a “Netscape Moment”?	2
AI and the Magical Moment	3
But...AI Is <i>Not</i> Magic	5
Moving Your Business from +AI to AI+	6
Before You Do Anything, Change Your Mental Model from +AI to AI+	7
The AI Ladder, Rebooted for GenAI	8
Before You Start Your Journey, Classify the Budget and Identify How AI Is	
Going to Help	11
Dimension One: Spend Money to Save Money, or Spend Money to Make	
Money? How Will AI Help Your Business?	11
Dimension Two: Categorize How the AI Helps Your Business	12
Use an Acumen Curve to Visualize How AI Helps Your Business	13
Where to Start? Here’s Our Helpful Advice	16
Become a Shifty Business: Shift Left, and Then, You Can Shift Right!	17
Every Day, We Walk by Problems That Can Be Solved or Made Better with	
Technology	18
Tips for Harnessing Foundation Models and GenAI for Your Business	24
Tip 1: Act with Urgency	24
Tip 2: Be an AI Value Creator, Not Just an Occasional AI User	24
Tip 3: One Model Will Not Rule Them All, So Make a Bet on Community	25
Tip 4: Run Everywhere, Efficiently	26
Tip 5: Be Responsible Because Trust Is the Ultimate License to Operate	27
And with That, Let’s Focus on the AI Part	27

2. Oh, to Be an AI Value Creator.	29
AI Through the Years: The AI “Time Lapse” Section	30
A Quick Bit on Foundation Models	31
Going a Little Deeper: The Evolution of Large Language Models and Comparing Supervised Learning with Self-Supervised Learning	37
AI Value Creation Should Be Your Destination	41
How Do You Consume AI: Be Ye a Value Creator or a Value User?	42
Planning Your AI Future: A Future with Many GenAI Models	45
It’s Time to Demystify and Apply AI	46
The Future of AI	53
Let’s Get into It	54
3. Equations for AI Persuasion.	55
Some Things Are Timeless	56
Tension Has Always Existed with Technology—Always	57
No Calculators Needed! Our Three Persuasion Equations	57
Equation 1: How to Grow GDP	60
Equation 2: What Makes for AI Success?	61
Equation 3: Find Your Balance—Navigate the Paradox	66
One Last Piece of Advice: See AI as a Value Generator, <i>Not</i> a Cost Center	71
Wrapping Up	72
4. The Use Case Chapter.	75
The Use Case Value Creation Curve	76
Going Horizontal Gets You the Most Vertical	78
Experimentation	78
Putting Your Data to Work	79
IT Automation	80
Code—The Language of Computers	83
Digital Labor and AI Assistants	90
Agents	93
The Business Lens: Use Cases—Horizontally Speaking	95
The Bonus (Horizontal) Use Case—Synthetic Data	95
A Smattering of Use Cases—Vertically Speaking	96
Agriculture	97
Accounting	98
Education	99
Healthcare	101
Insurance	102
Legal	104
Manufacturing and Production	106
Pharma	107

Endless Possibilities: More Industries Where GenAI Shines	109
The Building Blocks of AI	110
5. Live, Die, Buy, or Try—Much Will Be Decided by AI.	113
LLMs—The Stuff People Forget to Tell You	114
The Knowledge Cut-Off Date	115
LLMs Can Be Masters of Making It Up as They Go	115
Footprints in the Carbon: The Climate Cost of Your AI BFF	117
Copyright and Lawsuits	117
What About Digital Essence?	119
Your Expanding Surface Area of Attack	120
Data Privacy	123
Steal Now, Crack Later	124
Good Actor Levers for All Things AI	125
Fairness—Playing Fair in the Age of AI	126
Bias Here, Bias There, Data Bias Is Everywhere	127
Robustness—Ensuring Artificial Intelligence Is Unbreakable Intelligence	129
Explainability—Explain the Almost Unexplainable	133
Lineage—Tracing the Trail: Let Good Data Prevail	140
Regulations—The Section That Wasn’t Supposed to Be	141
What to Regulate—Our Point of View	142
Managing the AI Lifecycle	143
Wrapping It Up	145
6. Skills That Thrill.	147
Let the Skilling Begin	148
The Path to AI+ Requires Scaling Skills Across a Broad Spectrum of Roles	151
AI—Job Destroyer or Job Creator?	152
You’re Only Going to Get Checkmated if You Don’t Up Your Skills	152
Democratized Technology: The Job Creator	153
Levers of Clever: Unlocking a Skills Program That Lasts Forever	155
Lever 1: Start at the Beginning—Hire Employees Who Want to Know the “Why”	157
Lever 2: Recruit Digitally Minded Talent	159
Lever 3: Take Count—Inventory Your Skills	161
Lever 4: Plan for Everyone—A Plan Without Action Is a Speech	166
Lever 5: Embrace the Learning (and Forgetting) Curves	167
Lever 6: Combine Instruction + Imitation + Collaboration	169
Lever 7: Culture Matters—Be a Skills Verb, Not a Noun	173
Lever 8: Set the Organizational Tone for AI	174
Case Study: IBM’s Skills Challenge—the CEO Asked; We All Responded	175
The Final Word	178

7. Where This Technology Is Headed—One Model Will <i>Not</i> Rule Them All!.....	179
The Bigger the Better, Right? Perhaps at the Start, But That Was a Long Time Ago	180
The Rise of the Small Language Model	182
Data Curation Results in AI Salvation	183
Think About This When It Comes to Data Curation	190
Model Distillation—Using AI to Improve AI	190
Think About This When It Comes to Model Distillation	195
Where Are We Going Next? Small Language Models...Assemble!	197
Model Routing	197
Think About This When It Comes to Model Routing	202
Mixture of Experts (MoE) Architecture	203
Think About This When It Comes to MoEs	205
Agentic Systems	206
What’s Your Reaction to This Agent in Action?	208
A Little More on Agents	212
How Agents Are Built	213
Risks and Limitations of Agentic Systems	215
Three Tips to Get You Started: Our Agentic Best Practices	216
Think About This When It Comes to AI Agents	217
Wrapping It Up	218
8. Using Your Data as a Differentiator.....	219
Customizing Open Source for the Enterprise: A New Way of Looking at Enterprise Data	220
The Original Eras Tour: Looking Back a Few Decades on Data Representations	220
Stand Up and Represent!...Your Data	224
Step 1: It All Starts with Trust	224
Step 2: Representing your Enterprise Data within an LLM	228
Step 3: The Grand Finale: Deployment and Experimentation	238
The Future Is Open, Collaborative, and Customizable	239
9. Generative Computing—A New Style of Computing.....	241
The Building Blocks of Computing	243
Transformers—More Than Meets the AI	247
Not Back to the Future; Back to Computer Science	249
Doors Wide Open—Reimagining the Possible	251
How Models Are Built in Generative Computing	254
“Libraries” for Adding Capabilities to a Generative Computing System	255
The Quick Compare Summary—How You Use LLMs Today Versus Generative Computing	256

A Generative Computing Runtime—What Can We Program It to Do?	257
OpenAI’s Strawberry—A Berry Sweet Innovation	258
From Generative Computing to a Generative Computer—What Does All of This Mean for Hardware?	262
Experimenting with the Acceleration of AI at the IBM NorthPole	263
The Final Prompt: Wrapping It All Up	266
Index.....	267

Preface

Thrilling—the one word we use to describe the possibilities and eventualities that generative AI (GenAI) and agents will enable. From the boiler room to the board room, we think GenAI will truly impact every industry. As technologists, this level of excitement comes about only once, maybe every other decade, and this is why more and more, every day, the possibilities of GenAI are being recognized by people all over the world. Indeed, GenAI has the promise of a revolution, but this one will affect the high-status brainwork that the Industrial Revolution never touched.

As many know, Steve Jobs dropped out of college and went on to lead one of the most successful companies in the history of the world. Many know Jobs loved simplistic elegance and the beauty of things (like fonts on which he spent a great deal of time studying), and you can see that in Apple’s products to this very day. Many don’t know he was also fascinated with the efficiency of locomotion. He took particular interest in a study that looked at the least amount of energy a species would use to race one kilometer (0.62 miles). It might surprise you (it surprised us) to note that the winner in this category was a condor! Humans? About one-third down the list. But when humankind got on a bicycle, they blew everyone off the efficiency charts. He concluded that humankind could build tools (computers then, and now GenAI and agents) that can make us better—while other species must adapt (which takes a long time, if they can at all). In 1990, he likened a computer to a bicycle and noted, “What a computer is to me is the most remarkable tool that we have ever come up with. It’s the equivalent of a bicycle for our minds.” We’re pretty sure if he was around today, he’d note that if a computer is a bicycle for our minds, then GenAI and agents combine to become the bicycle for your business. And just as people have mastered the art of using Excel to manage their finances, track expenses, and create intricate color-coded charts (often with more complexity than necessary), these technologies are poised to become the standard tool for automating tasks, sparking creative ideas, and making it seem like you’ve worked tirelessly to achieve them.

While this all sounds promising, remember that *AI is not* a promise about prosperity, and a better world isn't guaranteed simply because you use GenAI or agents. Why? It can also have a dark side to it, which you will also learn about in this book.

Times Change, Technology Changes Faster

Let's get this out of the way: stuff is going to be out of date by the time you read this book. Today, writing about AI is like giving you stats on the number of images uploaded to the internet every second. We tried our best to keep stuff up-to-date (for example, DeepSeek came out, and we harassed our editor to open the manuscript to make some updates). And for sure, there will be new benchmarks, new papers, new state-of-the-art (SOTA) frontier models, GPUs, other accelerators, and more. We reference in this book how technology years used to be like dog years (1:7) and now they've become more like mouse years (1:30). So, bear with us when something new shows up that could be at odds with something in the book. *That said*, we think you'll find the main point of this book is to give you a mental model of the things to think about on your GenAI and agentic journeys, the things that deserve your attention, the questions to ask (yourselves and your vendors), and more. We think you'll agree that what we teach you in this book has a way longer shelf life than what GenAI model is today's SOTA and tomorrow's "Wait, people actually used that?" With that in mind, we encourage you to look beyond the stats or a model version and ensure you really absorb the advice that lies within—it comes from a place of success and failure, and a large corpus of real-world experiences and observations of what works and what doesn't.

For example, the buzz around DeepSeek in early 2025 (shortly before we went to print) thunderously demonstrated the very timeless points we are making in this book. Let us explain. The market mostly assumed that training cutting-edge models requires millions (or hundreds of millions) in investment with the latest, fastest, and greatest chips. That it had to be proprietary, and that trade secrecy was essential. DeepSeek proved otherwise, using the very things we've written about or mentioned in this book (Mixture of Experts, distillation, and reinforcement learning, among others) and some clever new optimizations (which we cover in this book too). It was released with a very permissive MIT license—that speaks to the open community we talk about in this book. So, while it may not show up in a benchmark comparison chart in this book, that was never the point because benchmarks are like a *Whac-A-Mole* game. But you will see the very concepts we lay out in this book go to work in the market and spur on new movements and innovations. And make no mistake about it: new innovations and techniques will arise, but we think those will tuck nicely into the playbook we're giving you in this book to put it to work.

GenAI Is a Lift, Shift, Rift, or Cliff Moment

We want you to think back to the first time you heard about GenAI. It's a phrase that really became part of the public conversation in, maybe, late 2022. We have seen new models, evolved models, and an explosion of open models. In a matter of months, GenAI transformed from an intriguing curiosity into a fundamental force driving business innovation, with a fresh wave of use cases and applications emerging each day.

There is such rapid growth that we can't predict exactly where we will all be 10 years from now—or even 10 months from when we finished writing this book. But one thing we're certain about: you're going to want to be actively engaged in shaping that journey—and hopefully that's why you're reading this book right now. We're at a moment in time here: one that is moving from a world of processes run by humans supported by technologies to one where processes will be run by technology that are supported by (or assisting) humans. This truly is a lift (good), shift (opportunity), rift (not so good), or cliff (what you're running toward if you don't upskill) moment for you as a leader, you as an individual, and for the companies you work for.

The future of AI is not one amazing model to do everything for everyone (you will hear us tell you time and time again in this book: *one model will not rule them all*). AI's future will not just be multimodal (seeing, hearing, writing, and so on); it will also most certainly be *multimodel* (in the same way cloud became hybrid). AI needs to be democratized—and that can only happen if we collectively leverage the energy and the transparency of open source and open science—this will give everyone a voice in what AI is, what it does, how it's used, and how it impacts society. It will ensure that the advancements in AI are not driven by the privileged few, but empowered by the many.

Indeed, the DeepSeek hoopla raises a bigger question: Who will shape the future of AI? Again, we think AI development cannot be controlled by a handful of players, especially when some may not share the same fundamental values, such as protection of enterprise data, privacy, transparency, and more. We can't let AI leadership slip to those with different values and priorities. That would mean ceding control of a technology that will reshape every industry and every part of society. And this is why we'll keep saying that innovation and true progress can only come by democratizing AI. We think that 2025 must be the year when generative AI gets unlocked from its confines within a few players; and into 2026, we hope that a broad swath of society won't just be using AI—many will be building it too.

As you read this book, you will see why we think a huge part of an enterprise's GenAI toolkit will be smaller open source models—this is how the future will be built. For too long, AI has been seen as a game of scale—where bigger models meant better outcomes. But the real breakthrough is as much about size as it is about efficiency. In our

work at IBM, we've seen that fit-for-purpose models have already led to up to thirty-fold reductions in AI inference costs, and made training more efficient and AI more accessible. There is no law of physics that dictates AI must remain expensive. The cost of training and inference isn't fixed—it is an engineering challenge to be solved. Businesses, both incumbents and upstarts, have the ingenuity to push these costs down and make AI more practical and widespread.

There's an old Chinese proverb about when is the best time to plant a tree. Whatever that time is (it varies depending on who tells you the story), it's in the past. But there is no argument about the next best time: today. We want to thank you for taking the initiative to read our book. We're hoping you'll be thanking us when you're done reading it because we're going to give you a framework so you can start making your GenAI and agentic plans and how you can effectively, safely, and responsibly put AI to work for business.

As You Journey into the Book

In this book, we are going to demystify AI—generative AI and agents. We'll explore a bit on how we got here, how it works, and many of the ways they are poised (and will) transform businesses and societies at unprecedented scale. We often refer to this point in time as a “Netscape moment” (Netscape being the world's first internet browser) because that's just how profound of an effect we think this technology will have on us all.

Before you dive in, here's a quick summary to give you the highlights. Consider it your cheat sheet—without the guilt and with all the good stuff. Use it as a trailer for the book ahead, or to jump into a section that really catches your attention.

Chapter 1, “AI to AI+: Generative AI and the ‘Netscape Moment’”

This is a business moment that rarely comes around; in fact, the last time we saw something this big, it was 1993 when a web browser (called Netscape) freed up the internet from the hands of the privileged few and democratized it for the many. Don't miss it. Things are going to change in the same way they did with the internet. In this moment, you ultimately won't compete against AI, but you will be competing against other companies using AI. Think about it. If you're a company spending 25% of your budget on customer service and another company shifts two-thirds of this same spend to have AI do most of the work...well. GenAI and agents will become a dividing line between which businesses will prosper, and which will struggle to keep pace. But always remember, AI is not magic. A couple of years from now, you will think back to this chapter and literally see which companies were the thrivers, which were the divers, and which ones showed up out of nowhere to become the new arrivers. Which one will you be?

Chapter 2, “Oh, to Be an AI Value Creator”

There are many ways to use AI. Be an AI Value Creator not just an AI User! Start your AI journey with the notion that your data is important and you shouldn't give it away. You'll need an AI platform to become an AI Value Creator. AI Value Creators accrue and create much more value than AI Users.

Chapter 3, “Equations for AI Persuasion”

We give you a productivity paradox. Today there are many factors working against business success (decreasing productivity, declining population rates, and more friction and cost when accessing debt). You have a unique opportunity to put AI to work against these forces, especially with AI-fueled productivity and digital labor.

Chapter 4, “The Use Case Chapter”

This is not about pet projects; this is about use cases that drive real value. When you step back, you start to realize that if you master the horizontal use cases of AI (the patterns and the things they can do, such as see, hear, analyze, and more), you will more masterfully choose the right vertical use cases for your business. Remember, computer vision is computer vision. Writing is writing. After all, to a computer, everything is just a bunch of numbers—even your Taylor Swift Spotify playlist. And don't forget, ensure you take your company over the Value Tipping Point.

Chapter 5, “Live, Die, Buy, or Try—Much Will Be Decided by AI”

AI that people trust is AI that people will use. You must decide up front if you are going to be a good actor or a bad actor—you've seen both around social media and other innovations. Make this decision up front! Why? The world needs regulation at the speed of right and too often, governments move at the speed of molasses. In the end, governments (and hopefully customers) are going to demand that your AI is both explainable and accountable. Stuff your AI journey backpack with fairness, robustness, explainability, and lineage—it's the ultimate gear for scaling new heights and taking in some truly breathtaking views. Ensure that these are forethoughts and not afterthoughts because, like we named in the chapter, “Live, die, buy, and try—much will get decided by AI.”

Chapter 6, “Skills That Thrill”

Because the half-life of technology skills is so short, know this: you will miss out on the amazing potential GenAI and agents can deliver to your business if you don't constantly upskill the many. Your teams need to know what AI can do, what it can't do, what to look out for, and more. No, you don't need everyone to have computer science degrees, but the only way to democratize AI for the many is to upskill the many. After all, how can you stop walking by problems every day that you could solve or make better with technology if you don't know that they are even fixable (using technology) in the first place?

Chapter 7, “Where This Technology Is Headed—One Model Will Not Rule Them All!”

Remember, one model won’t rule them all. A carpenter’s tool belt doesn’t have one tool; it has many tools. What’s more, smaller, more nimble models are showcasing incredible results while addressing some major challenges the world is facing with the traditional large language model (LLM) approach used by many today. But in the end, open access to multiple transparent and open models is going to give you the best chance at success.

Chapter 8, “Using Your Data as a Differentiator”

The title says it all: leverage your data as a differentiator. This pairs with being an AI Value Creator. When you step back and realize that perhaps 1% (at most) of enterprise data is in the commonplace LLMs you’re likely using today, you realize there is value to be had. Data is like a gym membership; if you don’t use it, you get nothing from it, but you also can’t just give it away.

Chapter 9, “Generative Computing—A New Style of Computing”

A glimpse into the future where generative AI takes its rightful place alongside classical computing and quantum computing as a new building block for applications called *generative computing*. This implies that the way we use LLMs is going to leverage software development methodologies that will broaden their applicability, safety, scale, performance, and more. We also expect (it’s happening today) our LLMs to reason more, take their time when warranted, and be thoughtful in their responses. This gives rise to a new area of optimization and “magic”—inference time. This change will pull forward new hardware and accelerators, creating a new compute stack, perhaps even a new generative computer.

We think—with the help of this book—when you look back at this moment in history, you will be able to do so fondly, as someone who embraced data as a true resource and used GenAI and agents as a utility to create value. Here we are—the start of your journey. Let’s get into it.

Conventions Used in This Book

The following typographical conventions are used in this book:

Italic

Indicates new terms, URLs, email addresses, filenames, and file extensions.

Constant width

Used for program listings, as well as within paragraphs to refer to program elements such as variable or function names, databases, data types, environment variables, statements, and keywords.

Constant width bold

Shows commands or other text that should be typed literally by the user.

Constant width italic

Shows text that should be replaced with user-supplied values or by values determined by context.



This element signifies a tip or suggestion.



This element signifies a general note.

O'Reilly Online Learning



For more than 40 years, *O'Reilly Media* has provided technology and business training, knowledge, and insight to help companies succeed.

Our unique network of experts and innovators share their knowledge and expertise through books, articles, and our online learning platform. O'Reilly's online learning platform gives you on-demand access to live training courses, in-depth learning paths, interactive coding environments, and a vast collection of text and video from O'Reilly and 200+ other publishers. For more information, visit <https://oreilly.com>.

How to Contact Us

Please address comments and questions concerning this book to the publisher:

O'Reilly Media, Inc.
1005 Gravenstein Highway North
Sebastopol, CA 95472
800-889-8969 (in the United States or Canada)
707-827-7019 (international or local)
707-829-0104 (fax)
support@oreilly.com
<https://oreilly.com/about/contact.html>

We have a web page for this book, where we list errata, examples, and any additional information. You can access this page at <https://oreil.ly/AI-Value-Creators>.

For news and information about our books and courses, visit <https://oreilly.com>.

Find us on LinkedIn: <https://linkedin.com/company/oreilly-media>.

Watch us on YouTube: <https://youtube.com/oreillymedia>.

The Collective Thank-Yous

Obviously, we have our personal list of people we need to thank, but a book like this doesn't come about because a whole bunch of other people didn't help us out in one way or another too. This is the section for those people—and we hope we get them all.

We'll start out with thanking Linda Snow. Linda read every page in this book, several times, asked clarifying and thought-provoking questions, showered us with unforgiving scowls when it came to some grammar issues, and her eye for even the smallest of details (we call her a comma queen), was a huge help to this book. A heartfelt thanks to you, Linda.

Rebecca Reyes leads skills at IBM, and she designed and cowrote **Chapter 6** with us. She helped flesh out a recipe that you can use to help move your organization's skills forward. As she's often heard saying to anyone who will listen, "One of the most important things you can brag about that impacts a business is your organization's learning culture." She is an incredible mind, leader, and human being...thanks, Rebecca!

Maya Murad and Anna Gutowska were a big help and pushed their obsession with AI agents on us from the very start: thank you! And thanks to João (Joe) Moura at CrewAI for some cool use case ideas and code samples.

Chapter 7 got a lot of help from Mikhail Yurochkin (and the rest of the model router team). Thanks for your great work, it was a privilege to showcase it in this chapter.

In **Chapter 8**, thank you to the entire InstructLab team, along with Kim Martineau (IBM Research Comms), Jeremy Eder (Red Hat), and Syeda Aameena Begum (an AI engineer in Bengaluru).

David Cox was a big part of **Chapter 9**—his insights into the future of GenAI and agents are nearly unparalleled—he is a natural-born storyteller.

Then there's the marketing team at IBM that works day in and day out communicating the value of the things we do. We thought we'd be remiss if we didn't call some of them out by name here because some of their work appears in some format (changed in tone and style) within this book. It starts with Sarah Benchaita, Sarah Meron, and Rebecca Neufeld. They are an incredibly talented and insightful bunch, and we are lucky to know them. We're also really thankful for the fabulous Lindsey Lurie. Others