



JUDITH HURWITZ
HENRY MORRIS
CANDACE SIDNER
DANIEL KIRSCH

AUGMENTED INTELLIGENCE

THE BUSINESS POWER OF
HUMAN-MACHINE COLLABORATION

 **CRC Press**
Taylor & Francis Group

AN AUERBACH BOOK

AUGMENTED INTELLIGENCE

THE BUSINESS POWER OF
HUMAN-MACHINE COLLABORATION

JUDITH HURWITZ
HENRY MORRIS
CANDACE SIDNER
DANIEL KIRSCH



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business
AN AUERBACH BOOK

Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2020 Hurwitz & Associates
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

Printed on acid-free paper

International Standard Book Number-13: 978-0-367-18489-6 (Hardback)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Contents

Dedications	<i>ix</i>
Contents	<i>xi</i>
Foreword	<i>xvii</i>
Preface	<i>xix</i>
Why This Book? Why Now?	<i>xix</i>
Why You Should Read This Book	<i>xxi</i>
What Is in This Book	<i>xxii</i>
About the Authors	<i>xxv</i>
Chapter 1: What Is Augmented Intelligence?	1
Introduction	1
Defining Augmented Intelligence	2
The Goal of Human–Machine Collaboration	3
How Augmented Intelligence Works in the Real World	3
Improving Traditional Applications with Machine Intelligence	4
Historical Perspective	5
The Three Principles of Augmented Intelligence	8
Explaining the Principles of Augmented Intelligence	8
Machine Intelligence Addresses Human Intelligence	
Limitations	9
Human Intelligence Should Provide Governance and Controls	11
Summary: How Augmented Intelligence and Artificial Intelligence Differ	12
	<i>xi</i>

Chapter 2: The Technology Infrastructure to Support Augmented Intelligence	15
Introduction	15
Beginning with Data Infrastructure	15
What a Difference the Cloud Makes	16
The Cloud Changes Everything	17
Big Data as Foundation	17
Understanding the Foundation of Big Data	18
Structured versus Unstructured Data	19
Machine Learning Techniques	20
Dealing with Constraints	20
Understanding Machine Learning	22
What Is Machine Learning?	22
Iterative Learning from Data	23
The Roles of Statistics and Data Mining in Machine Learning	24
Putting Machine Learning in Context	25
Approaches to Machine Learning	25
Supervised Learning	25
Unsupervised Learning	26
Reinforcement Learning	27
Neural Networks and Deep Learning	27
Evolving to Deep Learning	28
Preparing for Augmented Intelligence	28
Chapter 3: The Cycle of Data	29
Introduction	29
Knowledge Transfer	30
Personalization	30
Determining the Right Data for Building Models	31
The Phases of the Data Cycle	31
Data Acquisition	33
Identifying Data Already within the Organization	33
Reasons for Acquiring Additional Data	34
Data Preparation	35
Preparing Data for Machine Learning and AI	37
Data Exploration	38
Data Cleansing	38
Feature Engineering	41
Overfitting versus Underfitting	43

Overfitting versus Underfitting for a Model Predicting Housing Prices	45
From Model Development and Deployment Back to Data Acquisition and Preparation	46
Chapter 4: Building Models to Support Augmented Intelligence	47
Introduction	47
Explaining Machine Learning Models	48
Understanding the Role of ML Algorithms	48
Inspectable Algorithms	49
Black Box Algorithms	51
Supervised Algorithms	52
Creating a Gold Standard for Supervised Learning	52
K-Nearest Neighbors	53
Support Vector Machines	54
Unsupervised Algorithms	55
Understanding Reinforcement Learning and Neural Networks	56
The Value of Machine Learning Models	59
Summary	59
Chapter 5: Augmented Intelligence in a Business Process	61
Introduction	61
Defining the Business Process in Context with Augmented Intelligence	61
Weak Augmentation	62
Strong Augmentation	64
Strong Augmentation: Business Process Redesign	65
Augmented Intelligence in a Business Process about People	69
Strong Augmentation for Predictive Digital Marketing Campaign Management	69
Redefining Fashion Retailer Business Models with Augmented Intelligence	70
Business Model Changes at The Gap, Inc., Using Algorithmic Fashion Predictions	71
Another Fashion Retailing Business Model Using Algorithmic Predictions: Stitch Fix	72
Hybrid Augmentation	73
Summary	73

Chapter 6: Risks in Augmented Intelligence	75
Introduction	75
Providing Context and Understanding	76
The Human Factor	77
Understanding the Risks of a ML Model	78
The Importance of Digital Auditing	78
The Risks in Capturing More Data	79
Why It Is Hard to Manage Risk	80
Seven Key Risks	80
1. The Risk of Overfitting or Underfitting	80
2. Changing Business Processes Increases Risk	81
3. The Risk of Bias	81
4. The Risk of Over Relying on the Algorithm	83
5. The Risk of Lack of Explainability	84
6. The Risk of Revealing Confidential Information	84
7. The Risk of a Poorly Constructed Team	85
Summary	85
Chapter 7: The Criticality of Governance and Ethics in Augmented Intelligence	87
Introduction	87
Defining a Control Framework	88
Creating Your Augmented Intelligence Control Framework	91
Steps in Your AI Control Framework	92
Conducting a Risk Assessment	92
Creating Control Activities	92
Creating a Monitoring System	93
Data Privacy Controls	93
On an Organizational Approach to Controls	96
Summary	98
Chapter 8: The Business Case for Augmented Intelligence	101
Introduction	101
The Business Challenge	101
Taking Advantage of Disruption	103
Disrupting Business Models	103
Advantages of New Disruptive Models	104
Managing Complex Data	105

Creating a Hybrid Team	105
The Four Stages of Data Maturity	106
Building Business-Specific Solutions	110
Making Augmented Intelligence a Reality	110
How Augmented Intelligence Is Changing the Market	111
Summary	111
Chapter 9: Getting Started on Your Journey to Augmented Intelligence	113
Introduction	113
Defining the Business Problem	114
Establish a Data Culture	116
Moving Forward with the Foundation	117
Taking the First Steps	117
Selecting a Project that Can Be a Reference for Future Projects	118
Warning Signals	121
Summary	123
Chapter 10: Predicting the Future of Augmented Intelligence	125
Introduction	125
The Future of Governance and Compliance	126
Emergence of Different Jobs	126
Machines Will Learn to Train Humans	127
New Techniques for Identifying Bias in Data	128
Emerging Techniques for Understanding Unlabeled Data	129
Emerging Techniques for Training Data	129
Reinforcement Learning Will Gain Huge Momentum	130
New Algorithms Will Improve Accuracy	130
Distributed Data Models Will Protect Data	131
Explainability Will Become a Requirement	131
Linking Business Process to Machine Learning Models	132
Summary	132
References	133
Index	137



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Foreword

Artificial intelligence (AI) has been around for decades. AI has long piqued our imagination about what is possible with this powerful technology—for good and for evil. There are a vast number of instances in which AI has played a role in both our culture and in our desire to create more sophisticated systems. In fact, AI has captured the imagination of data scientists and fiction writers for decades. Alan Turing introduced seminal work on the meaning of artificial intelligence in the 1940s with his famous Turing test. In the 1960s project, ELIZA gave us a glimpse into the future of artificial intelligence–based communications. ELIZA, a project designed at MIT’s artificial intelligence lab, was one of the first natural language processing computer programs that created a way for humans to communicate with machines. Early science fiction novels were obsessed with artificial intelligence. For example, science fiction novels such as *Metropolis*, written by 1925 by German writer Thea von Harbou, imagined the world in the year 2026, where humans and robots would live together but engage in massive conflicts.

In the last few years, AI has re-emerged due to two main factors: the unprecedented compute ability to crunch huge amounts of data with improved machine learning algorithms, and the brilliant marketing buzz of AI and ML.

However, artificial intelligence is way more than a passing fad. It is already a major part of our daily lives and will continue to accelerate. It is clear that there is an enormous amount of accessible data. We now have the ability to synthesize this massive amount of data in order to better understand everything from changing customer behaviors to the ability to anticipate what customers will want to buy in the future. With the movement toward AI, we can improve our ability to more accurately predict future business trends and requirements. AI’s impact on modern medicine, for example, cannot be overlooked. AI will help to analyze complex data to support the physician’s ability to customize effective

treatments for patients. Across industries we are seeing important uses of AI that will be transformative.

Although AI has many types and forms, the most impactful in the immediate future is augmented intelligence. Augmented intelligence has profound ramifications because it empowers humans to truly leverage data to make better-informed decisions. The goal of augmented intelligence is not to replace people or automate them out of existence but to enable them to make better decisions based on complex data. In fact, the collaboration between humans and machines is the key enabler to our ability across industries to transition to the next stage of the information revolution.

One of the greatest impacts of augmented intelligence is the need to determine the accuracy of data in context. Understanding the context of data and the ability to actually trust this information is one of the most critical issues plaguing businesses today. If data can't be trusted or is inaccurate, decisions that are made based on that data could prove to be catastrophic. In addition to assuring data accuracy, data must be managed from a security, privacy, compliance, and risk perspective.

The book you are about to read is of great importance because we increasingly rely on machine learning and AI. Therefore, it is critical that we understand the ability to create an environment in which businesses can have the tools to understand data from a holistic perspective. What is imperative is to be able to make better decisions based on an understanding of the behavior and thinking of our customers so that we can take the best next action. This book provides a clear understanding of the impact of augmented intelligence on both society and business.

— Tsvi Gal, Managing Director
Enterprise Technology and Services
Morgan Stanley

Preface

Why This Book? Why Now?

Writing this book was a collaborative process among four seasoned professionals who have a common belief that when paired with human intuition and knowledge, artificial intelligence (AI) can change the world. The most pragmatic and useful way to benefit from AI and machine learning is to implement these powerful technologies as an augmentation to human intelligence. This hybrid approach—a partnership between humans and machines—is what is called *augmented intelligence*.

All four authors have a deep appreciation of the nuances of how we can harness the power of AI as a tool for transforming business. The value of our collaboration is that each of us brought a different perspective to a common goal of providing guidance and direction. Judith Hurwitz served as the author team leader. Judith has been a trusted advisor to many companies in a broad range of emerging technologies, ranging from data and analytics to cloud computing and business process. She has served on a variety of boards of advisors. In addition, Judith has coauthored 10 books, hundreds of e-books and articles, and is a frequent speaker and guest lecturer. Henry Morris is a technology thought leader with a PhD in philosophy from the University of Pennsylvania. Henry founded the analytics practice at the global research firm International Data Corporation (IDC), where he coined the industry term “analytic applications.” Henry has decades of expertise in analytics applications, business process management software, and the complete data life cycle. Henry has consulted with a variety of organizations and businesses across the globe. Candy Sidner is a renowned AI research scientist with a PhD in artificial intelligence from MIT. She has written more than 100 research papers and is a fellow of the Association for the Advancement of Artificial Intelligence. Candy’s involvement in the AI

community stretches from some of the earliest days of AI to today's emerging research areas. Daniel Kirsch is an attorney, senior consultant, thought leader, and author. Dan has written dozens of e-books and white papers on topics focused on the cloud, data, security, and compliance. Dan's research focuses on emerging technologies and their impact on businesses as well as the security, governance, and compliance implications of new technology.

The best part of writing this book was our weekly team meetings, where we discussed and debated the meaning of the movement to augmented intelligence and the need to create a hybrid approach that pairs humans with machines. Therefore, we bring a unique perspective to both the business and technical nuances required to use artificial intelligence on a collaborative platform that makes human responses to complex problems more achievable.

There is no doubt that artificial intelligence is powerful and provides huge benefits in being able to determine patterns that impact business outcomes. At the same time, the context and nuances of any field are vast. To make well-informed decisions, you must be able to assess data while utilizing human intuition and an overall understanding of the field in question. The human brain can make certain judgments and decisions incredibly quickly and accurately, but these choices are nearly impossible to codify into an artificial intelligence system. For example, you can't rely on a programmatic bot to suggest the right mix of painkillers to a patient without having a full understanding of that patient's physical and emotional state. You also need to understand the community the patient lives in and what support services are available. A person may appear on paper to be stable but might have some underlying issues in their home life that will impact treatment options. There is no machine learning model that can capture the full context of a person and their environment without human assistance.

As a team, we felt it was our obligation to put the extraordinary hype about artificial intelligence and machine learning in perspective. Venture capitalists have poured billions of dollars into companies that are promising to transform entire markets with intelligent systems that will understand everything and automate almost any task one can imagine. In addition, there is a wide variety of companies focused on AI, ranging from newly formed ventures to some of the most well-established public companies that are planning for their future. There are futurists who will tell you that artificial intelligence-based systems will be able to think at the same level as humans. We believe this claim is a misconception. There is no one technology that will have the power and intuition of human experts. The ability to translate complex data into applied knowledge may hold the key to solving some of the most complex problems we as humans face. It is our view that it is not enough to have a machine learning model and