MBR Management and Business Review

Special Double Issue: AI for Customer Engagement

Installation view of Refik Anadol: Unsupervised   **MoMA, New York, November 19, 2022**

Sponsored by business schools at Carnegie Mellon, Cornell, Dartmouth, UCLA, Baltimore, Michigan, Virginia, Vanderbilt, China Europe International Business School, City University of Hong Kong, Indian School of Business, INSEAD

Publisher's Letter to Company CEOs and Business School Deans

Greetings from the *Management and Business Review*! MBR, a peer-reviewed journal, is intended for executives, managers, students, and scholars. It is supported by over one hundred faculty members representing all of the top twenty business schools and more. It is sponsored by twelve leading business schools.

MBR's editorial team includes senior executives, more than one hundred professors from the top twenty schools, and the editors of many top research journals. Here are shareable links to this special double issue on AI for Customer Engagement as well as all three issues of MBR's 2021 debut volume:

https://mbrjournal.com/AI-Customer-Engagement-issue/

https://mbrjournal.com/2021-3-free-issues/

If you are reading the print issue, please visit mbrjournal.com to access these issues. This special issue explores how AI can increase customer engagement within the context of AI's impact on every aspect of our lives, work, and play. The debut volume of MBR includes articles from authors at all top-ranked schools and several Fortune-500 companies. Please visit our website, mbrjournal.com, to read *Forbes'* story on MBR.

To CEOs: Please share the links to the digital versions of this issue and the three issues of the 2021 debut volume with all your employees, customers, suppliers, and other contacts. Reading MBR articles is a highly effective way for your employees to keep up to date on the latest ideas in management from preeminent business scholars and leaders. You can subscribe to MBR for all your employees at a nominal price per employee. We would also be happy to help you customize MBR by adding special content tailored to your company's mission and the needs of your employees, customers, and suppliers. This customization will strengthen your relationship with your constituents and elevate your company's brand. Please contact coeditor in chief Kalyan Singhal at Ksinghal@ubalt.edu for details.

To Business School Deans: Please share the links to the digital versions of this issue and the three issues of the 2021 debut volume with all your alumni and other business and government constituents. We invite you to offer complimentary digital issues of MBR to your alumni. You will pay only a nominal price per alum. You also have the option of customizing MBR with content specific to your school. This customization will benefit your stakeholders and elevate your school's brand. Unlike some journals, which charge for classroom use, MBR makes its articles freely available for use in degree program classes at any school whose library subscribes to MBR. Please contact coeditor in chief Kalyan Singhal at Ksinghal@ubalt.edu to discuss pricing or customization options.

CONTENTS
AT A GLANCE

CONTENTS
AT A GLANCE

CONTENTS
AT A GLANCE

CONTENTS
AT A GLANCE

CONTENTS
AT A GLANCE

CONTENTS
AT A GLANCE

Advisors

- Regina Abrami, The Wharton School
- Ritu Agarwal, Johns Hopkins Carey Business School
- Raj Aggarwal, Kent State University Foundation
- Gad Allon, The Wharton School
- Edward Anderson, University of Texas at Austin
- Eugene Anderson, University of Pittsburgh
- Paul A. Argenti, Dartmouth College
- Linda Argote, Carnegie Mellon University
- Anil Arya, Ohio State University
- Baris Ata, University of Chicago
- Rob Austin, University of Western Ontario
- Karen Ayas, The Ripples Group
- Arun Balakrishnan, Formerly of Hindustan Petroleum
- Moloy Bannerjee, Bangalore
- Felix Barber, Ashridge Strategic Management Centre
- Richard Barker, Oxford University
- Caryn Beck-Dudley, AACSB
- CB Bhattacharya, University of Pittsburgh
- Amar Bhidé, Tufts University
- John Birge, University of Chicago
- Sam Bodily, University of Virginia
- Srinivas Bollapragada, General Electric
- Robert Bordley, University of Michigan
- Joseph Bower, Harvard University
- Tyson Browning, Texas Christian University
- Ryan Buell, Harvard University
- Andrew Campbell, Ashridge Strategic Management Centre
- Dennis Campbell, Harvard University
- Peter Cappelli, The Wharton School
- Jeff Cares, Alidade
- Glenn Carroll, Stanford University
- Bhaskar Chakravorti, Tufts University
- Richard B. Chase, University of Southern California
- Fangruo Chen, Shanghai Jiao Tong University
- Bruce Chew, Monitor Deloitte
- Sunil Chopra, Northwestern University
- Vidyanand Choudhary, University of California, Irvine
- Bhagwan Chowdhry, Indian School of Business
- Eric Clemons, The Wharton School
- Maxime Cohen, McGill University
- Morris Cohen, The Wharton School
- David Collis, Harvard University
- Charles Corbett, UCLA
- Stephanie Creary, The Wharton School
- Tugrul Daim, Portland State University
- Murray Dalziel, University of Baltimore
- Tom Davenport, Babson College
- George Day, The Wharton School
- David De Cremer, National University of Singapore
- Nicole DeHoratius, University of Chicago
- Marnik Dekimpe, Tilburg University & KU Leuven
- Brian Denton, University of Michigan
- Bert De Reyck, Singapore Management University
- Preyas Desai, Duke University
- Suzanne de Treville, University of Lausanne
- Robin Dillon-Merrill, Georgetown University
- Yu Ding, Texas A&M University
- Alexandre Dolgui, IMT Atlantique
- Matthew Drake, Duquesne University
- Jean-Pierre Dubé, University of Chicago
- Robert Eccles, Oxford University
- Andreas Eisingerich, Imperial College London
- Jehoshua Eliashberg, The Wharton School
- Omar El Sawy, University of Southern California
- Richard Ettenson, Thunderbird School of Global Management
- Paul W. Farris, University of Virginia
- Ken Favaro, act2
- Fred Feinberg, University of Michigan
- Kasra Ferdows, Georgetown University
- Charles Fine, MIT
- Baruch Fischhoff, Carnegie Mellon University
- Stewart Friedman, The Wharton School
- Alfonso Gambardella, Bocconi University
- Giovanni Gavetti, Dartmouth College
- Annabelle Gawer, University of Surrey
- Mary Gentile, University of Virginia
- Gerry George, Georgetown University
- Stanley Gershwin, MIT
- Pankaj Ghemawat, New York University
- Anindya Ghose, New York University
- Ranjan Ghosh, Indian Institute of Management, Ahmedabad
- Itay Goldstein, The Wharton School
- Mark Gottfredson, Bain & Company
- Vijay Govindarajan, Dartmouth College
- Linda Green, Columbia University
- Rajdeep Grewal, University of North Carolina at Chapel Hill
- Yael Grushka-Cockayne, University of Virginia
- Alok Gupta, University of Minnesota
- Sushil K. Gupta, Florida International University
- Stefan Haefliger, City University of London
- Andrei Hagiu, Boston University
- Nicholas Hall, Ohio State University
- Gary Hamel, London Business School
- Warren Hausman, Stanford University
- Constance Helfat, Dartmouth College
- Ola Henfridsson, University of Miami
- Teck-Hua Ho, National University of Singapore
- Andrew Hoffman, University of Michigan
- Kartik Hosanagar, The Wharton School
- Arnd Huchzermeier, WHU-Otto Beisheim School of Management
- Mark Huselid, Northeastern University
- J. Jeffrey Inman, University of Pittsburgh
- Ravi Jagannathan, Northwestern University
- Anjani Jain, Yale University
- Dipak Jain, China Europe International Business School
- Karuna Jain, Indian Institute of Technology, Bombay
- Mansour Javidan, Thunderbird School of Global Management
- Nitin Joglekar, Boston University
- M. Eric Johnson, Vanderbilt University
- Ajit Kambil, Deloitte
- Karthik Kannan, University of Arizona
- P.K. Kannan, University of Maryland
- Uday Karmarkar, UCLA
- Andrew Karolyi, Cornell University
- Sunder Kekre, Carnegie Mellon University
- Kevin Lane Keller, Dartmouth College
- L. Robin Keller, University of California, Irvine

- Pinar Keskinocak, Georgia Institute of Technology
- Mikko Ketokivi, IE Business School
- Wolf Ketter, Erasmus University
- Art Kleiner, Wise Advocate Enterprises and New York University
- Prabhudev Konana, University of Maryland
- Praveen Kopalle, Dartmouth College
- Ramayya Krishnan, Carnegie Mellon University
- Nirmalya Kumar, Singapore Management University
- Rajiv Kumar, Microsoft India
- V. Kumar, Indian School of Business
- Howard Kunreuther, The Wharton School
- Eva Labro, University of North Carolina at Chapel Hill
- Mary Lacity, University of Arkansas
- Karim Lakhani, Harvard University
- Hau Lee, Stanford University
- Josh Lerner, Harvard University
- Baruch Lev, New York University
- Arie Lewin, Duke University
- Marianne Lewis, University of Cincinnati
- Igor Linkov, US Army Engineer Research and Development Center
- Christoph Loch, University of Cambridge
- Yadong Luo, University of Miami
- John Paul MacDuffie, The Wharton School
- José A.D. Machuca, Universidad de Sevilla
- Costis Maglaras, Columbia University
- Joseph Mahoney, University of Illinois
- Ann Majchrzak, University of Southern California
- Manoj Malhotra, Case Western Reserve University
- Roger Martin, University of Toronto
- Carl Mela, Duke University
- Steven M. Miller, Singapore Management University
- Philip Mirvis, Global Network on Corporate Citizenship
- Sunil Mithas, University of South Florida
- Vikas Mittal, Rice University
- Antonio (Toni) Moreno, Harvard University
- Kara M. Morgan, Ohio State University
- Ram Mudambi, Temple University
- John Mullins, London Business School
- Haig Nalbantian, Mercer
- Satish Nambisan, Case Western Reserve University
- Paul Nunes, Accenture
- Vijay Padaki, The P&P Group
- Margherita Pagani, SKEMA Business School, France
- Geoffrey Parker, Dartmouth College
- Elisabeth Paté-Cornell, Stanford University
- Koen Pauwels, Northeastern University
- John Pearce, Villanova University
- Christine Pearson, Thunderbird School of Global Management
- Maury Peiperl, George Mason University
- Georgia Perakis, MIT
- Dino Petrarolo, Competitive Capabilities International
- Michael Pinedo, New York University
- Gary Pisano, Harvard University
- David Pyke, University of San Diego
- Arun Rai, Georgia State University
- Shivaram Rajgopal, Columbia University
- Jagmohan Singh Raju, The Wharton School
- Ananth Raman, Harvard University
- K. V. Ramanathan, University of Washington
- M. Rammohan Rao, Indian School of Business
- Michael Raynor, Monitor Deloitte
- Stefan Reichelstein, Stanford University
- Werner Reinartz, University of Cologne
- Ortwin Renn, University of Stuttgart
- John R. Roberts, University of New South Wales
- Qin Rongsheng, Beijing National Accounting Institute
- Johan Roos, Hult International Business School
- Aleda Roth, Clemson University
- Roland Rust, University of Maryland
- Hernan Saenz, Bain & Company and Cornell University
- Joseph Sarkis, Worcester Polytechnic Institute
- Garth Saloner, Stanford University
- Amy Salzhauer, Good Growth Capital
- Vallabh Sambamurthy, University of Wisconsin
- Nada Sanders, Northeastern University
- Sergei Savin, The Wharton School
- Mohanbir Sawhney, Northwestern University
- Alan Scheller-Wolf, Carnegie Mellon University
- Glen Schmidt, University of Utah
- Paul J. H. Schoemaker, The Wharton School
- Abraham Seidmann, University of Rochester
- Arijit Sengupta, Florida International University
- D. V. R. Seshadri, Indian School of Business
- Sridhar Seshadri, University of Illinois at Urbana-Champaign
- Suresh Sethi, University of Texas at Dallas
- J. George Shanthikumar, Purdue University
- Jason Shaw, Nanyang Technological University, Singapore
- Yossi Sheffi, MIT
- Max Shen, University of California, Berkeley
- Jagdish Sheth, Emory University
- Anil Shivdasani, University of North Carolina at Chapel Hill
- David Simchi-Levi, MIT
- Kingshuk Sinha, University of Minnesota
- J. Cole Smith, Syracuse University
- N. Craig Smith, INSEAD
- Edward A. Snyder, Yale University
- Deepak Somaya, University of Illinois at Urbana-Champaign
- Ron Soonieus, Boston Consulting Group and INSEAD
- Chester Spatt, Carnegie Mellon University
- Mark Spearman, Strategic Project Solutions
- Karen Spens, Hanken School of Economics
- James Spohrer, IBM
- E S Srinivas, Indian Institute of Management, Bangalore
- Kannan Srinivasan, Carnegie Mellon University
- Rajendra Srivastava, Indian School of Business
- Bradley Staats, University of North Carolina at Chapel Hill
- Christian Stadler, University of Warwick
- Martin K. Starr, Columbia University and Rollins College
- Jan-Benedict Steenkamp, University of North Carolina at Chapel Hill
- Fabian J. Sting, University of Rotterdam
- Donald Sull, MIT
- Jayashankar Swaminathan, University of North Carolina at Chapel Hill

Submissions

- Before submitting your paper, please review our editorial policies: https://mbrjournal.com/2020/02/17/editorial-policies/
- Please submit your manuscript at: https://mbrjournal.com/submit-manuscript-2/

ACKNOWLEDGMENTS

Cosponsored by the Association of National Advertisers' Global CMO Growth Council

MBR would like to thank
Nick Primola, Group EVP, ANA,
for his help and guidance.

INTRODUCTION

Jerry Wind, Margherita Pagani, and Jerry Dischler
Guest Editors of MBR's Winter & Spring 2023 Special Issue

At the end of November 2022 when OpenAI, a research lab in San Francisco, launched ChatGPT, its hugely popular chatbot, few people expected it to speed to more than 100 million users two months after the launch. By January it was widely viewed as the fastest-growing consumer app ever launched. For that reason, it seemed appropriate to begin this special issue of *Management and Business Review* on AI and customer engagement with an experiment.

We asked ChatGPT to introduce this issue, providing it with the following prompt:

"Given the growing importance of AI and its impact on every aspect of our lives, work, and play, explain how AI contributes to enhancing customer engagement. Follow the style of an academic journal such as *Management and Business Review.*"

ChatGPT's instant response was:

"Artificial Intelligence (AI) has become a key driver of innovation in many industries, and its impact on customer engagement is particularly noteworthy. Customer engagement refers to the extent to which customers are involved with, and invested in, a company, its products, and services. AI technologies can enhance customer engagement by enabling companies to provide more personalized, convenient, and efficient customer experiences.

"One way that AI contributes to enhancing customer engagement is through the use of chatbots. Chatbots are AI-powered virtual assistants that can communicate with customers through messaging apps, websites, or voice-activated devices. Chatbots can handle a wide range of customer service tasks, such as answering frequently asked questions, helping customers place orders, or resolving issues. Chatbots can also be programmed to respond to customers in natural language, making the interactions more human-like and intuitive. This can help improve customer satisfaction and foster a more positive brand perception.

"Another way that AI contributes to enhancing customer engagement is through the use of predictive analytics. Predictive analytics uses machine learning algorithms to analyze customer data and make predictions about future behavior. This can be useful for companies to better understand their customers and provide them with personalized experiences. For example, a retailer could use predictive analytics to make product recommendations based on a customer's purchase history, browsing behavior, and demographic information. By providing relevant and useful recommendations, companies can improve customer engagement and increase customer loyalty.

"Finally, AI can also enhance customer engagement through the use of virtual and augmented reality. Virtual and augmented reality technologies can be used to create immersive and interactive customer experiences that can enhance engagement and increase brand loyalty. For example, a fashion retailer could use virtual reality to allow customers to virtually try on clothing, or a car manufacturer could use augmented reality to allow customers to see what a car would look like in their driveway. These types of experiences can help customers better understand and connect with a brand, leading to increased engagement.

"In conclusion, AI is a powerful tool for enhancing customer engagement. Whether through the use of chatbots, predictive analytics, or virtual and augmented reality, AI technologies can help companies provide more personalized, convenient, and efficient customer experiences. By doing so, companies can improve customer satisfaction, increase customer loyalty, and ultimately drive business growth."

ChatGPT might not get a PhD for that explanation but, for a chatbot, it is an articulate and impressive response. It also sums up our motivation for producing this special issue of *Management and Business Review*.

To broaden our experiment beyond ChatGPT, we then presented the same prompt to Bing, which has integrated elements of AI in the latest version of its search engine, and to Google's Bard, which the company describes as "an experiment that lets you collaborate with generative AI."

Both Bing and Bard gave responses that were similar in some ways to that of ChatGPT. According to Bing, "AI contributes to enhancing customer engagement by providing personalized, interactive and value-creating service experiences that foster customer loyalty and satisfaction." Bing emphasized the importance of perceived interactivity, value co-creation, emotional intelligence and customer ability readiness. In each of these cases, the AI-powered bot explained the reasons for its answer. Bard's response to the same prompt was similar in that it focused

on how AI "contributes to customer engagement by personalizing experiences, analyzing customer data and identifying behavior patterns." When AI algorithms recommend products or services based on customer data analysis and identification of behavior patterns, they are "more likely to resonate with the customer," Bard said.

Still, there were some revealing differences between the three generative AI bots. Bing stated that "AI stimuli alone may not be sufficient to engage customers effectively. Some studies suggest that human interactions with service employees are still preferred by most customers and play a vital role in enhancing customer engagement. Therefore, a balance between AI and employee service may be optimal for engaging and retaining customers." Neither of the other two bots mentioned the need for human-AI collaboration or the limitations of AI. Bard noted: "By automating tasks, personalizing experiences, and providing insights into customer behavior, AI is helping businesses to create more engaging and rewarding experiences for their customers."

By bringing up patterns of customer behavior, Bard brought greater nuance to its response to the prompt.

But applications of AI in customer engagement go far beyond large language models or LLMs such as ChatGPT, Bing, or Bard. Leading businesses today are applying AI broadly in a number of domains and delivering impact at scale. The link between AI and customer engagement has, until now, rarely been examined in depth. In this special issue, we fill that void and explore various facets of the connection.

We, the editors of this special issue of *Management and Business Review*, have written this introduction to give our readers a snapshot of the issue. We ran a global competition to find the most insightful papers to include. The papers, arti-

cles, and book reviews that were selected are presented in five broad clusters, whose borders are fuzzy. These groups represent the different ways in which AI can help build customer engagement: personalization, automation, prediction, generating customer insights, and enabling omnichannel engagement. A brief description of each paper, grouped by cluster, may be found in the Contents at a Glance at the beginning of the issue.

Scenarios for the Future

This special issue also features a report of a workshop hosted by Meta in New York City. *Management and Business Review* collaborated with the Global CMO Growth Council to organize this forum, which took place on July 26, 2022. Its purpose was to discuss three possible scenarios for the future, identify the most likely one, and determine what steps companies should take today to prepare. It relied on a scenario-planning exercise in which leading marketing experts from business and academia came together to predict the likely course of AI and customer engagement over the next five years. They did so by examining the most optimistic, pessimistic, and realistic scenarios for the future. Knowing that AI is making ever deeper inroads into every aspect of business and society, transforming finance, health care, manufacturing, and also marketing, participants considered, how AI will transform customer engagement by 2027.

The forum was held in 2022, before the tidal wave of generative AI crashed upon the world's attention. It is clear that generative AI programs such as ChatGPT, Bing, Bard, and GPT-4 will have a transformational impact upon every industry. The forum's most likely scenario and its implications are still valid, but recent advances in generative AI will act as a booster, turbocharging the adoption of AI and its use in customer engagement. This makes the need to act even more urgent. Ignoring generative AI is not an option.

In addition to considering the three 2027 scenarios, it is therefore essential to look further ahead, by at least two decades. So at the end of this issue is a twenty-year view based on Kai-fu Lee and Chen Qiufan's book, titled *AI 2041: Ten Visions for Our Future*. Kai-Fu Lee is the CEO of Sinovation Ventures and the author of *AI Superpowers*, a book that deals with the emerging rivalry in AI between the U.S. and China. A former president of Google China, he is now co-chair of the Artificial Intelligence Council at the World Economic Forum. Chen Qiufan, also known as Stanley Chan, is an author and translator and the founder of the content studio Thema Mundi. He was previously Kai-Fu Lee's colleague at Google. He is the president of the World Chinese Science Fiction Association and the author of *Waste Tide*, published in English in 2019.

The book *AI 2041* is an unusual blend of science fiction and analysis. It describes ten scenarios or visions that are likely to become reality over the next twenty years. It consists of ten science fiction stories, each set in a different part of the world, that are written by Stanley Chan, with commentary by Kai-Fu Lee. The combination of compelling science fiction with real-life analysis makes it fascinating reading.

A Call to Action

Now that you, our readers, know what to expect from this special issue of *Management and Business Review*, how can you make sure that you get the most from it? It is our hope that the papers, articles, and book reviews that you read here will motivate and inspire you about AI as it redefines the way in which organizations engage with their customers and other audiences. We hope that you will turn inspiration into action, using what you have learned to construct your own experiments, wherever you may be, and using AI to build greater engagement with your own customers. If that happens, we will all have succeeded. It is an exciting journey. Bon voyage! ■

©shutterstock.com/thodonal88

The Mastercard Digital Engine™:
Using AI to Spot Micro Trends for Effective Customer Engagement

Raja Rajamannar
Mastercard

Introductory Note

Raja Rajamannar, chief marketing and communications officer of Mastercard, describes the dwindling of consumer attention and explains how marketers can follow Mastercard's example and use emerging AI technology to engage consumers meaningfully, authentically, and in real time.

C onsumer attention spans are getting shorter by the day, and for good reason. Every day, people are bombarded with roughly 10,000 advertising messages. Seeking a reprieve, many consumers look for platforms that provide what I call a 'pure ad-free heaven' devoid of advertising interruptions.

According to a 2020 industry study by Edelman, nearly seven in ten adults worldwide use one or more methods to avoid ads, changing their media habits to see fewer ads (49 percent) and using ad blockers (48 percent).[1] Those numbers will surely rise in the years ahead.

This noisy ad landscape is a nightmare for marketers who want to reach consumers with messages they would actually welcome – relevant, personalized offers that drive engagement and boost the advertiser's return on investment (ROI).

Moreover, a 2021 Gartner report found that 63 percent of digital marketing leaders struggle to deliver personalization because they are still scaling up their use of emerging technologies.[2] Consumers are ignoring most branded content or deeming it irrelevant – costing brands a fortune in lost customers and revenue.

It is clear that the old, traditional ways of advertising are no longer viable.

It is clear that the old, traditional ways of advertising are no longer viable. Marketers must learn how to engage consumers through new, alternative routes that leverage emerging technologies with fresh strategic thinking and an updated roadmap.

That is the mindset which led digital marketers at Mastercard to reimagine their framework. From the onset, the goal has been to engage consumers with content that is personalized, contextually relevant, and in the moment. To accomplish this, the team created the Mastercard Digital Engine™ with artificial intelligence (AI) and machine learning (ML) at its core.

The AI-Powered Mastercard Digital Engine™

How does it work?

1. The engine spots micro trends all over the world as it wades through billions of conversations on the Internet. Micro trends range from a new cuisine like sushi burritos and disappointment over ballet dancers going on strike, to a rise in contactless payment options at the onset of the pandemic.

2. Since Mastercard creates experiences for consumers by tapping into their passions – think culinary, travel, sports, music, entertainment, and more – the engine instantly matches micro trends with our current experiences and offers, suiting them to consumers' preferences and passions.

3. Armed with a deep analysis of the micro trend, the marketer can then decide whether to engage consumers and launch a campaign with a Mastercard Priceless experience, offer, or reward. The campaign can be launched on multiple media platforms in minutes (not months) with custom content pulled from Mastercard's rich library of relevant copy. And it can be taken down the second the trend has run its course.

Not only is AI facilitating the Mastercard Digital Engine™ and enabling it to act with zero or near-zero lag for maximum impact, ML and natural language processing (NLP) are also integral to accelerating its impact, allowing us to measure the effects in real time, optimizing who receives what content and when. This trifecta is game changing.

The AI processes we have employed thus far include: named entity recognition (NER) algorithmic approach, graph-based (PKE), and unsupervised keyword extraction (YAKE), to name a few. The algorithm is then further refined by using Word Mover's Distance, Isolation Forest, and One Class SVM to remove irrelevant micro trends.

In addition to its real time response that matches consumer interests and desired experiences with a micro trend, the engine also boosts the relevance of our content by capturing both short-lived and long-term trends to help marketers deliver a personalized message in a contextually adaptable way.

The following examples showcase these capabilities.

Micro Trend No. 1: Celebrity's Breaking News

A celebrity announced a big career transition, and the news generated significant online buzz. The Mastercard Digital Engine™ spots a spike in conversation about the celebrity and matches it with a behind the scenes video on Priceless, Mastercard's consumer platform which features experiences, offers, and rewards for cardholders. A creative campaign is instantly created from personalized, contextually relevant and timely content.

In an authentic and meaningful way, we were able to drive higher engagement and click-through rates. The campaign, which was active for a two-day period, drove significantly greater awareness, consideration, and conversion than campaigns delivered by traditional methods:

- 100 percent higher engagement rates (engagement rates: social platforms calculate this as interactions divided

by a number of impressions; interactions are likes, shares, comments, etc.)

- 254 percent higher click through rates as compared to benchmarks
- 85 percent reduction in cost per click

The relevance of the message, an assertive creative campaign, and the real time model of the engine allowed us to generate a cost-effective outcome that surpassed our objective's key performance indicators (KPIs).

Micro Trend No. 2: Latest Tourist Passions

A European tourism board enlisted the Mastercard Digital Engine™ to help promote its country to residents of an adjacent nation, hoping to drive cross-border travel and thereby generate economic growth. We worked closely with the tourism board to define the strategy and implement the tactics.

The engine spotted relevant micro trends and matched them with compelling content and offers available to consumers instantly. The resulting ads fell into three passion categories – travel, culinary, and culture – driving clicks to the tourism board's website.

This AI-powered campaign delivered higher metrics than similar but traditional campaigns that the tourism board ran, targeting the neighboring nation. Both were on the same media platforms and had similar goals. Mastercard Digital Engine™ delivered:

- 16 percent lower cost per reach
- 20 percent more people reached
- 87 percent lower cost per engagement
- 25 percent higher rate of engagement
- 38 percent lower cost per click
- 96 percent higher click-through rates

Micro Trend No. 3: Trendy Traveler Experiences

A national airline teamed up with Mastercard to build a sweepstakes campaign designed to attract travelers to a culinary series in a popular domestic destination. The Mastercard Digital Engine™ spotted seven micro trends, matching them to pre-identified digital ads, called creatives, and copy in various categories. We engaged consumers with customized content that was contextually relevant and in the moment. The campaign generated the following results compared to campaigns that employed traditional methods:

- 29 percent lower cost per click (CPC)
- 37 percent higher click-through rate (CTR)
- 32 percent lower cost per engagement (CPE)
- 43 percent higher engagement rate (ER)

Key Findings

We know the engine works because over the last two years Mastercard has launched more than 500 successful campaigns across twenty countries for either itself or its partners. Our initiatives have covered a wide variety of categories ranging from sports and music to travel and food, as well as e-commerce, philanthropy, and diversity, equity, and inclusion.

AI is a one-of-a-kind technology that is driving a higher standard of marketing effectiveness. A game changer, the Mastercard Digital Engine™ delivers much higher campaign metrics than those of traditional campaigns. We have measured a statistically significant subset of these campaigns, applying strict test and control methodology to arrive at the following results:

- **Reach multiple**
  - average is 1.8 times
  - median is 2.0 times
  - range is 0.6-9.1 times

- **Click-through rates multiple**
  - average is 4.1 times
  - median is 2.2 times
  - range is 0.8-21.4 times
- **Engagement rates multiple**
  - average is 3.2 times
  - median is 2.0 times
  - range is 0.8-15.9 times

Sentiment, which we have also observed, usually varies across regions. In the Latin American market where the engine is used quite broadly, we have observed an improvement of eight percentage points in positive and neutral sentiment toward Mastercard.

We are also seeing differences across passion categories. For instance, categories with larger followings, such as sports, music, and food, do well. Localization is also important. Soccer does well across most regions but does not fare well in the U.S.

Another key discovery concerns the duration of campaigns. Short campaigns with a small number of creatives generally do not perform as well as longer campaigns or campaigns with more creatives.

We have always known that, without the power of AI, this could not be done in anything near real time.

Nothing Will Influence The Marketing Field More Than AI

We started this journey a few years ago with the goal of designing a marketing tool that could successfully cut through today's crowded and cluttered landscape, credibly reaching consumers who tune out messages because they prefer frictionless experiences. Fueled by AI, the Mastercard Digital Engine™ is an effective and efficient machine that builds campaigns that drive real time, impactful engagement and ROI.

Turning to the power of AI was a natural step from the beginning. From learning deeply about consumers to enabling hyper-personalization, to optimizing programs, we have always known that, without the power of AI, this could not be done in anything near real time.

Since the engine's inception, we have been delighted to see how quickly AI technology and ML are advancing. ML ensures that the engine improves with every execution, picking up better, more dynamic trends by the day. Thanks to ML, we will continue to hone our ability to spot appropriate micro trends. We also look forward to advancements in empirical measurement.

The progress of NLP is also fascinating to observe and put into use. Scaling our engine globally and into many languages would have been hugely more difficult even a couple of years ago. Think of the colloquial nuances and adaptation of each language in different locales; the Spanish spoken in Spain is quite different from the Spanish spoken in Mexico or Puerto Rico.

AI, ML, and NLP allow us to act with zero or near-zero lag time for maximum impact. And that impact can also be measured in real time so campaign optimization can happen instantly. (We define optimization as determining who receives what content, and when.)

If you still believe that AI's speed can be matched by manual methods, think again. The lines between machines and humans are blurring in ways that were once unimaginable. As marketing practitioners aided by AI, we can have a finger on the pulse of every stage of the marketing life cycle, make sense of it, and act, yielding highly effective outcomes instantly, an effect termed *quantum marketing.*

If you still believe that AI's speed can be matched by manual methods, think again.

The Mastercard Digital Engine™ is engaging, enlightening, and empowering. Because of its globality, we are seeing many common threads and unifying themes. We are also seeing notable differences in how it operates across regions. For example, the development of the natural language models that exist today varies dramatically from one language to the next. We need to bring more access and equality into the information technology framework, ensuring inclusion and equity. As we continue to improve our capabilities, this social responsibility must be a bigger part of a digital marketer's handbook.

Our engine has taught us so much. Principally, it has ingrained in us an unswerving belief that AI is a necessity for engaging customers in a meaningful and authentic way in real time. It is the only alternative to manual methods today because making a human connection is everything. ■

Author Bio
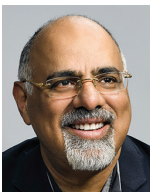
Raja Rajamannar is Chief Marketing and Communications Officer and President of Healthcare at Mastercard. He is leading the company's innovative marketing transformation and evolving its identity. Raja is President of the World Federation of Advertisers and serves on several boards. He is an inaugural member of Forbes' CMO Hall of Fame and one of Business Insider's World's Most Innovative CMOs. Raja is the author of the Wall Street Journal best-selling book *Quantum Marketing*.

Endnotes

1. https://www.insiderintelligence.com/insights/ad-blocking/
2. https://www.gartner.com/en/newsroom/press-releases/-gartner-says-63--of-digital-marketing-leaders-still-struggle-wi

The Alibaba Challenge:
How to Effectively Engage with a Billion Customers

©shutterstock.com/Oleksandra Bezverkha

Yitong Wang
Alibaba

Ofer Mintz
University of Technology
Sydney and Tel Aviv University

Depin Chen
Alibaba

Kehan Chen
Alibaba

Alibaba conducts millions of transactions involving billions of dollars each year. Yitong Wang, Ofer Mintz, Depin Chen, and Kehan Chen describe how the company uses AI-powered chatbots to complement its human resources so as to maximize engagement with its millions of customers and with its vendors.

Alibaba Group, China's largest e-commerce company and one of the world's biggest companies, has nearly one billion annual active Chinese consumers who daily participate in hundreds of millions of transactions through its Taobao e-commerce platform.

During Alibaba's busiest shopping period in 2021 – Taobao's Double 11 shopping festival – its platforms facilitated 540 billion RMB (US $85 billion) in transactions over a two-day period. The sheer number of customer interactions in Taobao transactions makes it difficult for Alibaba to keep up with proactively engaging with its customers.

Further, Alibaba needs to resolve several million customer service queries each day. The queries can come from end-user consumers or business merchants in Alibaba's two-sided platform. Consumers may have questions about the Alibaba platforms or be unsatisfied with products they purchased from merchants on the platform. Merchants could have questions about end-user consumers or about the platforms.

Taken together, Alibaba faces resource restrictions that limit its ability to engage with customers solely by human service interactions. Since 2015, Alibaba's response has been to implement artificial intelligence (AI) chatbots, supplemented by human service interactions, to both proactively and reactively engage with its customers.

Alibaba uses AI chatbots to handle customer engagement for more than two million daily sessions and over 10 million lines of daily conversation.

Today, Alibaba uses AI chatbots to handle customer engagement for more than two million daily sessions and over ten million lines of daily conversation on Taobao's two-sided platform, representing about 75 percent of Alibaba's online and 40 percent of phone hotline consultations.

Not only has the use of AI chatbots raised customer satisfaction by 25 percent, based on initial results, it has saved the company more than one billion RMB annually (~US $150 million) by employing AI instead of human contact center agents.

Alibaba's AI chatbots

Alibaba employs five AI chatbots to cater to the heterogenous demands from customers on Taobao's two-sided platform.

- The Wanxiang-bot assists merchants on the Taobao platform, performing tasks such as resolving questions about the platform's rules, activities, and service issues.
- The Alibee Shop bot assists merchants with end-user consumer interactions on Taobao, helping with service issues and direct engagements between merchant and consumers.
- The Alime bot helps the end-use consumer. It is employed in online and phone hotline channels and relies on a rich set of interactive user interface components that can provide text dialogues, cards, graphics, videos, and other conversational interactions between robots and consumers.[1] Alime bot also possesses duplex voice dialogue

Table 1: Pros and Cons of Using AI or Human Agents for Customer Engagements			
	AI	**Human**	**Examples**
Pros	• Faster and more accurate responses for cases with clearly defined questions and answers • Easily scalable, low marginal cost • High supply elasticity	• More accurate responses for complex and non-standardized cases • Useful when customers need to detail evidence of problems • Better understanding of long-tail problems not well defined by algorithmic rules	• Lead time: For Taobao's Double 11 shopping festival, human services typically need three months for hiring, preparation, and training in comparison to AI services that require only a one-day lead time focusing primarily on increasing cloud computing capacity. • Costs: AI-based customer service costs less than 1 percent of human customer service.
Cons	• Unable to handle complex scenarios with many rounds of conversation • Strong dependence on algorithms and training data • Weak emotional connection with customers	• Slower responses • Quality is highly dependent on personnel • Low supply elasticity • High cost	• Interpretability: AI-based bots need to be provided with at least hundreds of training samples; humans can read the question and (mostly) understand the problem • Number of interactions per customer: The current AI-based bot service is limited to two rounds of customer dialogue; humans can continue for longer customer dialogues

Figure 1. Alibaba's Five AI-bots Employed for Customer Service

Figure 1. Alibaba's Five AI-bots Employed for Customer Service

(arrow direction indicates which party initiates the conversations)
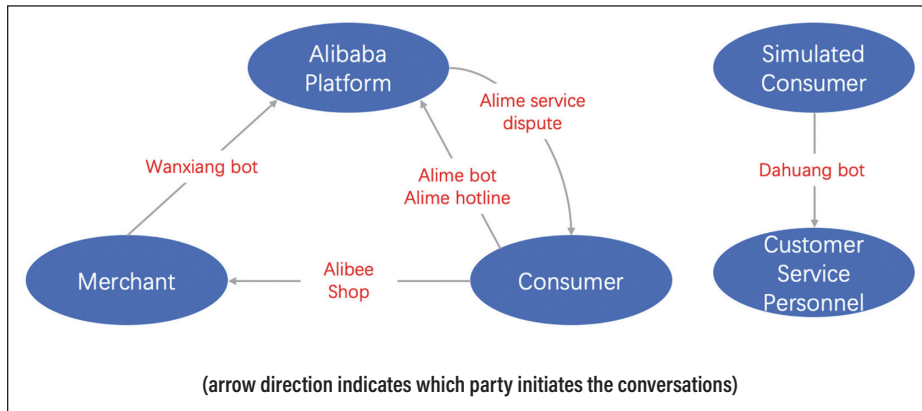
capabilities and a targeted user interface to serve users who prefer to consult by telephone.

- The AI bot complements Alime by proactively engaging with end-user consumers and acting as an intermediary during service disputes with business merchants. The AI bot's underlying algorithms use transaction and conversational information to evaluate service disputes and make automated judgments. The AI bot then calls end-users to discuss the dispute decision. If consumers are satisfied with the AI bot's decision, the case is closed; if not, the AI bot helps consumers fill out appeal forms to have the case resolved by a human evaluator.

- The Dahuang bot helps train customer service personnel by simulating a larger and more diverse set of consumer and merchant encounters than a non-AI training system. (See Figure 1.)

Implementation challenges and solutions

The first challenge Alibaba faced was organizational hesitancy: engineers and leaders were wary of AI's ability to positively engage with and resolve consumer and merchant queries. To overcome this hesitancy, the firm implemented a fast-fail strategy[2] that contin-

> The first challenge Alibaba faced was organizational hesitancy: engineers and leaders were wary of AI's ability to positively engage with and resolve consumer and merchant queries.

ually tested the AI against human responses in small real-world experiments.

The pilot proved that the AI chatbots outperformed their human counterparts, delivering superior customer satisfaction scores and improving merchants' first-contact resolution scores. Those results led the organization to fully buy in to implementing AI customer engagement services.

The second challenge involved two technical issues. Many initial customer queries involved a user-intent classification problem: Queries were stated similarly but had different underlying intents. For example, "I need help with my order" can mean needing help with tracking an order, getting a refund, or many other intents. In addition, many customer queries involved a long-tail problem: extremely low likelihood of certain types of queries. Since niche topics by nature gener-

ate less data, queries about them tend to lead to less accurate chatbot answers – which lowers customer satisfaction.

To overcome the first technical problem, Alibaba developed an extensible multitask learning paradigm using a meta lifelong learning framework that learned robust text representations across tasks and employed a least recently used (LRU) replacement policy to manage model deployment and memory resources.[3] To overcome the latter problem, Alibaba implemented a multi-grained interactive matching network for few-shot text classification that leverages a dynamic routing algorithm in meta-learning to better adapt and generalize unseen classes while also providing more memory-based flexibility.[4]

AI vs. human results

Overall, Alibaba's AI-based customer satisfaction scores exceeded or matched human interactions in most product categories. Crucially, it provided customers much quicker responses and enabled Alibaba to engage with customers at all hours of the day.

A/B testing results reveal that, in its first few weeks, Alibaba's proactive AI-based intermediary dispute service resulted in a 25 percent increase in customer satisfaction over its previous non-AI-based dispute resolution procedures.

Meanwhile, Alibaba's training bot is assisting more than 1,500 customer service personnel daily. In addition, it is reducing customer service personnel training time by more than 20 percent.

Finally, Alibaba has realized cost savings of over a billion RMB annually (~US $150 million) by employing AI over humans for service engagements.

Lessons learned and next steps

Alibaba employs its five AI-based chatbots in more than eighty

Table 2: Summary of Alibaba's Five AI Bots

AI-bot	Alibee Shop	Wanxiang-bot	Alime (bot & phone hotline)	Alime (customer service dispute)	Dahuang-bot
Target Customer	End-user consumers	Business merchants	End-user consumers	End-user consumers	Customer service personnel
Year Service Began	2018	2015	2015 (online) 2018 (hotline)	2021	2017
Purpose	Assist consumers with merchant-related questions on Taobao platform (via chat)	Assist merchants with help about the Taobao platform (via chat)	Assist end-user consumers use of the Taobao platform (via chat and voice hotline)	Assist end-user consumers service disputes with business merchants (via chat and voice hotline)	Training customer service personnel with a broader and more diverse range of questions than human-based training (via chat and voice hotline)
Examples	Service engagements involving direct merchant-to-consumer engagements	Service engagements involving business merchant questions on the platform's rules, activities, and service issues	Service engagements involving customer service questions and customer engagements	Service engagements resolving dispute resolutions	Training customer service personnel on how to respond to customer service engagements
Key Technical Features	· QA match · Intention recognition · Rec. system	· QA match · Intention recognition	· QA match · Intention recognition (bot) · Visual QA (bot) · Automated speech recognition (hotline) · Multi-turn dialogue (hotline) · Text-to-speech (hotline)		· QA match · Intention recognition

Question-Answer (QA) match: the need to match the user's query to the title of the corresponding knowledge base. The user's description is based on natural language, which has rich and diverse expressions for the same problem. Only AI models based on deep learning can provide highly accurate and normalized solutions based on the quantity of knowledge data need to address the user's questions.

Multi-turn dialogue: the ability for an AI bot to engage in back-to-back conversational turns with a user.

Alibaba-related apps, including Taobao, Xianyu, Fliggy, Hema, and Lazada. It took Alibaba developers one year to create the initial AI bot. However, this process has become much more efficient, and it now takes Alibaba's developers only one month to develop a new AI bot for its platforms, although they have to continually improve each bot.

Alibaba has also learned that customers have accepted the chatbots and appreciate their role in resolving many, though not all, types of customer service queries. It has learned that it's critical to constantly test the structure to see what types of help customers are willing to accept from AI chatbots.

A key lesson Alibaba learned was that, despite its chatbots'

success, AI cannot and should not completely replace human customer service agents. But each should be deployed in the scenarios that best suit their abilities.

For simple FAQ questions, AI can directly reply to users. However, for complex complaints and disputes, AI can attend to labor-intensive tasks, such as collecting informa-

Alibaba has realized a cost savings of over a billion RMB annually (US $155 million) by employing AI over humans for service engagements.

tion on appeals and vouchers, and possibly make initial decisions, but humans will still need to review the information, as well as related non-standardized materials, to make final judgments.

Going forward, Alibaba will continually invest significant resources in making a seamless human-machine collaboration. For example, Alibaba's AI bots continuously monitor whether customers encounter obstacles and whether the AI-based service can understand and resolve their queries.

Whenever appropriate, the AI bot automatically transfers customers to human service agents and then prompts and provides those agents with essential information to avoid asking the

customers to repeatedly describe the problem.

At the same time, the AI bot enhances customer interactions with human agents by providing data-informed solutions and particular phrasing recommendations, as well as issuing immediate warnings to human agents if they behave improperly.

Finally, Alibaba believes that, although technologically proficient AI natural language processing models and training data are important, establishing an organizational mindset which offers customer-centric AI-based solutions to efficiently solve customer problems is more important than just implementing AI.

So, Alibaba is continuing to refine its AI bots through hundreds of releases, iterations, and improvements over time, with the ultimate goal of solving customers' problems in the most efficient and satisfying method possible, whether through AI or human service. ■

Author Bios
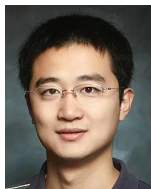
Yitong Wang is head of the Data Science and Business Insights Department in the Customer Experience Business Group at the Alibaba Group. Yitong earned his undergraduate and master's degrees from Tsinghua University and his PhD from the University of California, Irvine. His research has been published in the Journal of Consumer Psychology, Strategic Management Journal, and Journal of Experimental Psychology: General, among others.

Ofer Mintz is an Associate Professor of Marketing at the University of Technology Sydney, Visiting Associate Professor of Marketing at Tel Aviv University, and author of *The Post-Pandemic Business Playbook: Customer-Centric Solutions to Help Your Firm Grow*. His research focuses on marketing metrics, analytics, strategies, and marketing's role in start-up firms. His research has been published in the *Journal of Marketing, Marketing Science*, and *World Economic Forum*, among other managerial and academic outlets.

Depin Chen is the Director of Applied Algorithms in the Customer Experience Business Group at the Alibaba Group. Depin is an expert in developing recommender systems, conversational AI, anomaly detection, and more. He earned his Bachelor's degree and PhD the University of Science and Technology of China, in 2005 and 2010, respectively. He has several publications in leading data and computer science conference proceedings such as ICDM, SIGIR, and PAKDD.

Kehan Chen is Manager of the Chatbot Algorithm Team in the Customer Experience Business Group leading the AliMe Conversation AI algorithm team at the Alibaba Group. Kehan's research interests include natural language understanding, human-computer interaction, and dialogue systems. Kehan earned his master's degree from Zhejiang University, Hangzhou, China and has authored or coauthored several papers in leading data and computer science conference proceedings, including at Interspeech and KDD.

Endnotes

1. Alibaba's AI-bot's interactive UI components are open sourced and available for others to employ (see https://chatui.io/).

2. For further details on fast-fail strategies, please see: Mintz, Ofer, and Eric Knight (2022), "Beyond Usual: [Six] Ways Leading Firms diverge from Business as Usual." *Management and Business Review*, forthcoming.

3. For further details, please see: Wang, Chengyu, Haojie Pan, Yuan Liu, Kehan Chen, Minghui Qiu, Wei Zhou, Jun Huang, Haiqing Chen, Wei Lin, and Deng Cai (2021), "MeLL: Large-scale Extensible User Intent Classification for Dialogue Systems with Meta Lifelong Learning," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, New York, NY, USA: Association for Computing Machinery, 3649–3659.

4. For further details, please see: Geng, Ruiying, Binhua Li, Yongbin Li, Jian Sun, and Xiaodan Zhu (2020), "Dynamic Memory Induction Networks for Few-Shot Text Classification," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Online: Association for Computational Linguistics, 1087–1094.

AI in Personalized
Product Recommendations

Kartik Hosanagar
The Wharton School, University of Pennsylvania

Dokyun Lee
Questrom School of Business, Boston University

Consumers today are presented with a vast wealth of product offerings in digital marketplaces and storefronts. It is essential that managers consider how to ensure that shoppers find their company's products. Kartik Hosanagar and Dokyun Lee explain the power of recommender systems, fueled by AI, to do just that.

G iven the explosion in product offerings in digital marketplaces and storefronts, managers must consider how consumers will find products of interest among seemingly endless alternatives. Recommender systems are an important solution to the problem. These systems combine data drawn from sources including clickstream, purchases, product ratings, user profiles, and social networks to predict which products are best suited to a particular customer. They help consumers to become aware of new products and to select desirable ones from a myriad of choices. For firms, recommender systems have the potential to convert browsers into buyers, cross-sell products, and increase customer loyalty. Introduced nearly two decades ago, recommender systems are becoming ever more relevant and their impact ever stronger as available data and product assortments increase.

Recommender systems exert a significant influence over consumer choice.

Recommenders are known to exert a significant influence over consumer choice. According to a McKinsey & Co report, 75 percent of viewing hours streamed on Netflix and nearly 35 percent of sales at Amazon originate in automated recommendations.[1] In 2016, Netflix estimated that their recommendation engine is worth $1 billion or more per year.[2] In e-commerce, recommender systems have been conservatively estimated to drive a 15 percent increase in product views and a 7.5 percent increase in total conversion, from click to purchase.[3] Firms such as TikTok have also found personalization technologies to be a source of competitive advantage.[4]

Recommender Systems Design Practices

The design of recommender systems is an active area of research, with many articles available on its specific facets.[5] More generally, recommender designs can be classified as either *content-based* or *collaborative filter-based* systems. Content-based systems use either product metadata (e.g., author, genre, musical attributes) or deep-learning-processed raw content (e.g., soundwave data for songs) to recommend items similar to those that a user rated highly. Collaborative filters, in contrast, recommend what similar customers bought or liked. For example, the classical customers "who bought item X also bought item Y."

Let us consider the design of three music recommendation services: Pandora, Last.fm, and Spotify. These services use three different approaches to recommender system design.[6]

Pandora's online radio service emerged from the Music Genome Project, a research effort that uses musical attributes to describe songs. Musicologists listened to tracks and assigned over 450 attributes to each. These range from the extent of instrumentation in the music to more esoteric attributes such as "rhythmic syncopation." Once a user indicates that they like a song on Pandora, the algorithm finds other songs that have similar musical attributes. For example, if a user chooses "Thunder" by Imagine Dragons, Pandora might recommend "Ride," a pop song by Twenty One Pilots, which features reggae beats. Pandora indicates that the song was recommended because it features "a dub production, a reggae feel, acoustic rhythm piano, use of a string ensemble, and major key tonality."

Pandora's approach depends on having the detailed attributes of every product (in this case, musical tracks), and is therefore called a content-based recommender system. Collecting this data manually is incredibly time-consuming and expensive. Moreover, knowing these attributes isn't all that useful when a retailer sells a wide range of products. If you've only ever bought books on Amazon.com, your taste for post-war thrillers will do little to help the company know what music to recommend, let alone which shoes, couches, or cars.

Collaborative filtering is an alternative approach that you might recognize through the "People who bought this also bought…" and "People like you also liked…" recommendations that we often see on Amazon and other websites.[7] The collaborative filtering approach used by Last.fm identifies the users who listen to the song "Thunder" and notes the other songs they have been listening to. This approach does not depend on rich metadata and as a result is unlikely to come with detailed domain-specific explanations like those provided by Pandora.

What collaborative filters lack in depth of knowledge, they make up for in simplicity. They are easy to implement and roll out in a short time. They have therefore become the most popular class of automated product recommenders on the Internet. They work very well in practice and the social pull of knowing what others are listening to or buying likely helps to drive consumers' choices.

These two design approaches can be used to recommend not only music but also videos (Netflix, YouTube), news (Google News) or, really, any product on retail websites (Amazon).

The advantage of a collaborative filtering design is that there is no need for detailed metadata

and it is better able to capture the social benefits of shared consumption. However, the approach has its drawbacks. One major drawback is that collaborative filters create a sales concentration bias whereby popular products are more likely to be recommended. Because they recommend products based on what others have consumed, collaborative filters cannot recommend products that are not yet popular even if they would have been rated favorably by the consumer.[8] These systems also cannot explain the reasons for their recommendations, since the algorithm has no knowledge of product characteristics. Finally, collaborative filters also have a cold-start problem with new products; they are unlikely to recommend new products because they have no previous customer purchase or ratings data. These three drawbacks are well addressed by content-based designs. But content-based recommenders are difficult to build because they require detailed metadata which may not be available.

Because the two design approaches have different strengths and weaknesses, hybrid designs that combine the simplicity of collaborative filters with the impartiality of content-based designs are highly appealing. But how do we extract detailed metadata about products in a way that does not involve significant manual effort? To avoid the time and effort involved in having musicologists listen to and collect musical attributes for each and every track, we can use AI to automate the work. Spotify's hybrid design does precisely that.[9]

Spotify crawls the web to look at blog posts and other online discussions about music, figuring out the kind of descriptive language that listeners use to discuss different songs and artists. These terms then become attributes of the

songs. But new or niche songs aren't discussed as much online, so the data this process finds about such songs is insufficient. Spotify compensates for such deficits by using a machine learning algorithms to analyze each song and extract audio characteristics such as tempo, loudness, key, and tonality.[10] Algorithmic approaches that take raw audio data and find interesting patterns to be used as attributes in recommender systems are called feature engineering algorithms. At the end of this feature engineering, Spotify's algorithms have both Last.fm's understanding of music listening patterns and Pandora's deep understanding of the music itself.

The end result is Spotify's Discover Weekly, an algorithmically generated personalized weekly playlist. According to the company, as many as 8,000 artists get over half their streams from users listening to their Discover Weekly playlists.

Recent advances in deep neural networks, a type of machine learning algorithm, now enable automated feature engineering, or feature learning. Beyond just the audio data as mastered by Spotify, deep neural networks can natively handle and process multi-modal data, that is any combination of both structured and unstructured data like text, photo, video, audio, to automatically engineer features and utilize content-based recommendations for any products including images, videos, people, and even firms.[11]

But as with any algorithmic solution, the broader impact of recommender systems on consumer choice needs close attention. It is worth noting that while some of the aforementioned design tradeoffs, such as specific explanations of content or ease of cold starting, are well understood, we know very little about

how these different designs affect consumer engagement. This could be a fertile area of research. Specifically, while recent research has shown the positive impact of recommenders on short-term consumer engagement, including product views, purchases, and consumption diversity, the long-term impact of recommenders on consumer engagement and retention calls for more investigation.[12] Furthermore, different types of recommenders have been shown to influence consumption diversity differently. While collaborative filters have been shown to decrease aggregate consumption diversity, hybrid systems do not seem to have the same affect.[13] This difference is important because a recent study from Spotify provides correlational support linking the diversity of consumption to long-term user engagement, including conversion and retention.[14]

Lee and Hosanagar found that recommenders work with other e-commerce features, such as review ratings and descriptions, to influence consumers' purchase decisions. In addition, Karim et al found that focusing only on recommendation accuracy at the expense of other objectives results in the recommendation of harmful content and even negative affects on mental health.[15] Poorly designed recommendation systems, especially in social news feeds, can also create filter bubbles in which consumers are exposed to a narrow set of content instead of a range of perspectives, fueling the fragmentation of society.[16] However, Hosanagar et al analyzed the impact of a content-based music recommender design and did not find evidence that it fragments users. Instead, they found that the design increases users' overlap in consumption. Still, over-representation of subgroups of customers with specific taste could introduce

bias into the recommender system, perpetuating feedback loops in the system and giving rise to unfair distribution of attention. For example, the system might only recommend items preferred by the over-represented group at the expense of other subpopulations within the system.[17] This bias can then marginalize or shift the preferences of the affected subpopulations. This erosion occurs because the algorithm's parameters are estimated from the initial starting data, which then influences recommendations with both selection bias and a preference-shifting effect, in which the user's preference is changed simply because the recommendations are made.[18] Deldjoo et al. provides a survey of research on fair recommender systems.[19]

These issues suggest that while recent applications of machine learning address some of the technical roadblocks to recommender design, a broader social science perspective is urgently needed. Incorporating such a perspective will recognize that recommender algorithms cannot be evaluated merely on the accuracy of their recommendations, the proportion of recommended products that users find relevant. Instead, managers should conceptualize

customer satisfaction or engagement more holistically, incorporating the recommendations' relevance with their impact on consumer well-being (avoid recommending content which is potentially addictive or otherwise harmful to mental health), social fragmentation (tending to narrow users' perspectives), breadth of consumption (tending to expose users to new topics and product categories), and long-term customer retention.

While recent applications of machine learning address some of the technical roadblocks to recommender design, a broader social science perspective is urgently needed.

Conclusions

Personalized product recommendations help consumers to discover products of interest and sort through the myriad of choices available online. They also help firms to improve customer retention and to cross-sell and upsell products. Despite their value, early recommender designs have many

limitations, ranging from the need of content-based recommenders for rich product metadata to collaborative filters' popularity bias, cold-start problems, and inability to explain the details of their recommendations. In recent years, machine learning has been applied to create hybrid systems that combine the best of both approaches, as Spotify has done. The method shows great promise in terms of incorporating new kinds of unstructured data to generate personalized recommendations. With the number of people online continuously increasing[20] and the rapid approach of an AI-augmented creator economy and metaverse, we expect recommender systems to have an even greater impact on consumer choices and engagement in the coming years.[21] However, we caution managers to not just focus on short-term customer engagement metrics, but also to monitor and evaluate long-term customer engagement and the societal impacts of large-scale personalization. Recent advances in recommender systems have demonstrated that it is not necessary to sacrifice the accuracy of recommendations in order to increase their diversity, which could circumvent the filter bubble effect without causing a drop in customer engagement.[22] ◼

Author Bios

Kartik Hosanagar is the John C. Hower Professor of Technology and Digital Business and a Professor of Marketing at The Wharton School. Kartik's research focuses on the digital economy, particularly the impact of analytics and algorithms on consumers and society, Internet media, Internet marketing, and e-commerce.

Kartik was named one of the world's top forty business professors under forty, and has received many teaching awards. He has consulted for Google, American Express, Citi, and more.

Dokyun "DK" Lee is a Kelli Questrom Chair Associate Professor of Information Systems Management and Computing and Data Sciences at Boston University. He studies the responsible application, development, and impact of AI in e-commerce and the digital economy, focusing on the economic impact of textual data and content extraction, understanding, and engineering. His research, which has won several awards, is supported by Adobe, Bosch Institute, Google Cloud, Marketing Science Institute, McKinsey & Co, Nvidia, and Net Institute.

Endnotes

1. MacKenzie, Ian, Chris Meyer, and Steve Noble. "How retailers can keep up with consumers." McKinsey & Company 18 (2013): 1.

2. https://www.businessinsider.com/netflix-recommendation-engine-worth-1-billion-per-year-2016-6

3. Lee, Dokyun, and Kartik Hosanagar. "How do product attributes and reviews moderate the impact of recommender systems through purchase stages?." Management Science 67.1 (2021): 524-546.

4. Wall Street Journal, "Inside TikTok's Highly Secretive Algorithm." July 21, 2021 https://www.wsj.com/video/series/inside-tiktoks-highly-secretive-algorithm/investigation-how-tiktok-algorithm-figures-out-your-deepest-desires/6C0C2040-FF25-4827-8528-2BD6612E3796

5. For a review of early design approaches, see: Adomavicius, Gediminas, and Alexander Tuzhilin. "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions." IEEE transactions on knowledge and data engineering 17.6 (2005): 734-749. A review of more modern deep-learning approaches is available in: and Zhang, Shuai, et al. "Deep learning based recommender system: A survey and new perspectives." ACM Computing Surveys (CSUR) 52.1 (2019): 1-38. To understand the potential problems of the latest recommender models, see: Ferrari Dacrema, Maurizio, Paolo Cremonesi, and Dietmar Jannach. "Are we really making much progress? A worrying analysis of recent neural recommendation approaches." Proceedings of the 13th ACM conference on recommender systems. 2019. Meanwhile, an examination of the latest trends and frontiers in recommender systems may be found in: Jannach, Dietmar, et al. "Recommender systems: Trends and frontiers." AI Magazine (2022).

6. Hosanagar, Kartik. "A Human's Guide to Machine Intelligence: How Algorithms Are Shaping Our Lives and How We Can Stay in Control." New York: Penguin Viking (2019).

7. Breese, John S., David Heckerman, and Carl Kadie. "Empirical analysis of predictive algorithms for collaborative filtering." arXiv preprint arXiv:1301.7363 (2013).

8. Fleder, Daniel, and Kartik Hosanagar. "Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity." Management Science, Vol. 55, No. 5, pp. 697-712 (2009). and Lee, Dokyun, and Kartik Hosanagar. "How do recommender systems affect sales diversity? A cross-category investigation via randomized field experiment." Information Systems Research 30.1 (2019): 239-259.

9. http://benanne.github.io/2014/08/05/spotify-cnns.html

10. https://hackernoon.com/spotifys-discover-weekly-how-machine-learning-finds-your-new-music-19a41ab76efe (october 10, 2017)

11. LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. "Deep learning." nature 521, no. 7553 (2015): 436-444.

12. Kumar, Anuj and Kartik Hosanagar. "Measuring the Value of Recommendation Links on Product Demand." *Information Systems Research*, vol. 30, No. 3, pages 819-838 (2019).

13. Hosanagar, Kartik, Daniel Fleder, Dokyun Lee, and Andreas Buja. "Will the global village fracture into tribes? Recommender systems and their effects on consumer fragmentation." Management Science 60, no. 4 (2014): 805-823.

14. Anderson, Ashton, et al. "Algorithmic effects on the diversity of consumption on Spotify." Proceedings of The Web Conference 2020. 2020.

15. Karim, Fazida, Azeezat A. Oyewande, Lamis F. Abdalla, Reem Chaudhry Ehsanullah, and Safeera Khan. "Social media use and its connection to mental health: a systematic review." Cureus 12, no. 6 (2020).

16. Pariser, Eli. The filter bubble: How the new personalized web is changing what we read and how we think. Penguin, 2011.

17. Chen, Jiawei, et al. "Bias and debias in recommender system: A survey and future directions." arXiv preprint arXiv:2010.03240 (2020).

18. Adomavicius, Gediminas, Jesse C. Bockstedt, Shawn P. Curley, and Jingjing Zhang. "Do recommender systems manipulate consumer preferences? A study of anchoring effects." Information Systems Research 24, no. 4 (2013): 956-975.

19. Deldjoo, Yashar, et al. "A Survey of Research on Fair Recommender Systems." arXiv preprint arXiv:2205.11127 (2022).

20. Statista, "Global digital population", April, 2022, https://www.statista.com/statistics/617136/digital-population-worldwide/

21. Facebook IQ, "Digital Research and Insights How AI Will Elevate the Future of Creativity" October 26, 2021, https://www.facebook.com/business/news/insights/how-ai-will-elevate-the-future-of-creativity?ref=search_new_1

22. Kunaver, Matevž, and Tomaž Požrl. "Diversity in recommender systems–A survey." Knowledge-based systems 123 (2017): 154-162.

©shutterstock.com/elenabsl

Hyper-Personalization for Customer Engagement with Artificial Intelligence

Thomas H. Davenport
Babson College and Oxford Business School

Personalization based on customer attributes and behavior is a familiar concept among marketers, and artificial intelligence is making it increasingly effective. AI-based hyper-personalization employs both sophisticated methods and far more data than previous methods and is far more precise as a result. Thomas H. Davenport discusses the role of AI in personalization as well as the growing backlash against personalization fueled by data privacy concerns.

Personalization is one of the most important ways marketers can use data, analytics, and artificial intelligence (AI) to increase customer engagement. It suggests to potential customers that the offered product or service is particularly suited to their specific needs and desires. It promises to value customers' attention and time, bringing to their attention only offerings consistent with their interests.

Personalization—also known as one-to-one (or 1:1) marketing—is not new. The concept has its roots in the idea of segmentation, in which marketers treat different groups differently. Personalization is ostensibly about appealing to specific individuals, but has traditionally relied upon placing people in groups such as gender, sociodemographic status, geographical residence, age, etc. Generally speaking, the more attributes marketers use to segment customers, the closer segmentation comes to personalization. However, until recently, few firms had the data or modeling ability to actually create a different offer for each consumer. Full personalization was more a concept

than a reality in the first decade or two of its existence; in practice it was not distinguished from segmentation.

Both segmentation and one-to-one marketing were primarily focused on increasing sales rather than improving customer engagement with a product, brand, or company. Customer engagement generally refers to developing an ongoing relationship and a broader customer experience, rather than a particular sales transaction.[1] One popular measure of engagement proposed by consultants Bain & Co., the Net Promoter Score (NPS), has been tied to company-level shareholder return growth through repeat and referred sales revenues.[2] However, many academic studies have found little or no relationship between NPS and consumer-level sales.[3] The idea that there is a relationship between personalization, engagement, and sales is thus rooted only in logic, rather than deep empirical results.

Over the last decade, with the rise of e-commerce and online marketing, personalization has been used to target digital advertising, offers, and other customer-oriented content. The vast amount of data involved in digital customer relationships, and advances in AI methods, make it increasingly possible to approach the ideal of a unique offer for each customer.

However, most approaches to personalization have not been very sophisticated or effective. They have usually fallen well short of the 1:1 ideal because they lacked accurate and detailed data for personalization, or because they used relatively primitive methodological approaches and algorithms. Today, when sufficient data is available, some methods can help marketers to realize the promise of true personalization, here termed *hyper-personalization*.

Many managers lack a good means of characterizing the technology and sophistication of alternative personalization approaches. Vendors tout their own approaches without mentioning limitations. Each approach has its own data requirements and technical underpinnings.

There are also privacy tradeoffs involved in personalization. The potential value of personalization is highly appealing to marketers and product or service designers, but we must address important issues concerning customer privacy and perceptions of invasiveness. Personalization based on data should therefore generally be conducted with transparency and customer permission.

What to Personalize?

Personalization is applicable to a wide variety of marketing approaches. Perhaps the most common is personalized advertising, either for digital or analog (direct mail, television, etc.) ads.

> The potential value of personalization is highly appealing to marketers and product or service designers, but we must address important issues concerning customer privacy and perceptions of invasiveness.

Offers and discounts on particular products and services can also be personalized. In online commerce, search rankings and other aspects of web and mobile sites can be personalized. These personalization approaches are typically intended to increase conversions, from initial click to final purchase, rather than overall engagement.

Marketers today can easily personalize digital content, whether in marketing messages, web pages, or other forms of expression. Many assume that it builds customer engagement more than transaction-oriented digital advertising. And personalized content has moved beyond business. Political campaigns are increasingly personalized, as are news feeds, which has given rise to concerns about the filter bubble effect, in which people only see content from sources with which they agree.

We can also personalize the attributes of products and services. Financial services firms personalize investment portfolios with AI advice bots. Clothing companies offer shoes and other goods which can be personalized at manufacture. Even medical companies increasingly offer services and treatments that are targeted to a patient's genome or other biological makeup, termed *precision medicine* or *personalized medicine*.

Many companies now offer personalization in behavioral recommendations. Startups employ *precision nudges* to encourage weight loss, good economic habits, and overall health and nutrition.[4] Because these personalized services benefit customers beyond encouraging them to purchase an individual product or service, they do seem to advance customer engagement.

Personalization has often been viewed as relating to offers—ads, promotions, new product announcements, and the like. A personalized *next best offer* is intended to motivate only the customer's next purchase. Firms aiming to increase customer engagement are more partial to the term *next best action*. AI systems designed for this broader concept may advise customers to purchase an additional product or service, but they may also recommend

actions or content that advance the customer relationship without making a new sale.

Recommended next best actions may include content about how the customer can improve their experience with the product or service, how the product or service can be used most effectively, or simply how to live a better life. Whenever possible, managers should ensure that this advice be personalized not only with regard to the product or service purchased, but also according to other customer attributes that increase the likelihood of their relevance.

Morgan Stanley's wealth management team, for example, has created a next best action system that provides machine learning-based personalized investment recommendations to its customers (mediated through investment advisors, who send the recommendations to customers). Investment advisors who use the system significantly outperform those who do not in growing the assets under management, and have higher productivity and frequency of engagement with customers.[5]

Benefits of Personalization

Personalization has many potential benefits. Survey research attests to some of the reputed benefits, while others are suggested by online consumer behavior. Most involve increased sales or conversion rates rather than broader metrics of customer engagement. A McKinsey 2021 survey, for example, found that 71 percent of consumers say they expect businesses to recognize them and to personalize product or service offers to their interests. Seventy-six percent are frustrated by an absence of personalization. In addition, companies surveyed by McKinsey that say they personalize report higher levels of revenue growth than those who do not.[6] Other less recent surveys also show a strong consumer preference for personalization.

However, the landscape surrounding the tradeoff between personalization and privacy is changing. Two 2021 KPMG surveys of businesses and consumers depict the challenge companies face when collecting data on consumers, which is necessary for effective personalization. Among 250 business executives, 70 percent reported that they expanded their collection of personal consumer data during the previous year. However, 86 percent of consumers said they have growing concerns about data privacy, and 30 percent said there are no circumstances under which they would share data with businesses. Only 12 percent said they would share data to personalize ads, and 17 percent would do so to help companies improve their products and services.[7]

As customers become more sensitive to data privacy, they may be less interested in personalization, or at least may require a higher degree of accuracy in personalizing to justify the tradeoff. It remains difficult to establish the exact frontier at which the desirability of personalization is outweighed by the desire for personal privacy, in part because the two traits are hazy in the minds of many consumers.

Contemporary AI-Based Approaches to Personalization

Firms have employed AI, in the form of rule-based systems, for personalization for decades, but until recently it has been a relatively blunt instrument. The most sophisticated, precise, and difficult form of personalization—in other words, hyper-personalization—requires machine learning models. Unlike rule-based personalization, machine learning can employ multiple different customer, product, or contextual attributes with few complications, while tending to be much more precise than rule-based approaches. It allows a company to produce millions of unique offers, drawn from many different groupings or segments. Machine learning-based personalization can even approach the elusive 1:1 segmentation to which marketers have long aspired. This approach has been theoretically possible for several decades, but lacked the necessary data. For example, the UK-based supermarket chain Tesco pioneered hyper-personalization in the late 1990s and early 2000s using its Clubcard loyalty program data to generate 12 million unique offers for grocery promotions. More recently, in the inflationary economy of 2022, Tesco focused on personalizing discounts for Clubcard holders.[9]

Hyper-personalization approaches that use machine learning can be centered on the item, the customer, or a combination of the two. Item-centric approaches like collaborative filtering rely on the "people who bought this item also bought this other item" approach, generating coefficients of item relatedness based on purchase data. All a company needs to know to use this approach is what items customers have expressed interest in or bought.

Some companies go deeper into the item-centric approach by classifying the attributes of their products. Netflix's well-known recommendation engine, for example, classifies each movie and TV show by multiple product attributes, such as subject, stars, directors, and the like. It can then recommend content to customers which has the same attributes as content they've purchased.

In addition to information about titles, Netflix's personalized recommendations are based on self-reported customer entertainment preferences, viewing history, and ratings of watched titles, as well as other members with similar tastes and preferences (it categorizes

more than 2000 different taste communities), and situational factors such as time of day, device type, and length of a typical viewing session. The engine uses the information to present a personalized set of titles to each viewer.[10] Netflix's personalization has clearly contributed to a better customer experience, and led to high subscriber growth for many years (though it has declined somewhat recently due to economic and post-pandemic contraction) and to higher combined usage than cable and satellite viewing combined.[11]

Customer-centric approaches to personalization might employ not only past purchases, but also customer demographic data, recent life events, estimated income levels, communications channel preferences, and responses to previous offers. They combine these variables to develop a predictive model of how a customer will respond to the personalized offer. They then use the model to score each potential customer in terms of the likelihood of purchasing a product or category. By deploying many different models, a company could approach making a different offer to each customer. The grocery store chain Kroger, for example, has 60 million customers in its loyalty programs and delivered 1.9 billion personalized offers to them in 2021, using a large-scale machine learning model developed by 84.51°, its data and analysis subsidiary.[12]

These offers are primarily focused on encouraging sales transactions rather than on customer engagement. Yet 84.51° and Kroger are also beginning to emphasize personalized nutritional information and recommendations that may boost engagement over time. The retailer's OptUP program uses Kroger loyalty card data to calculate a nutrition score from a customer's recent purchases. Shoppers can also browse an app while shopping to see nutrition scores of individual prod-

ucts and receive "Better for You" recommendations of similar but healthier products.

While companies like Kroger have voluminous data on shopper behavior, what usually makes customer-centric models difficult is obtaining the necessary data. To employ them, a company needs extensive data on customer attributes and labeled outcome data, such as whether customers purchased the product or category, or responded to an offer. Loyalty programs allow companies to track many aspects of customer behavior over time. In placing digital ads, companies have traditionally used cookies as an excellent source of data about websites that customers have visited, which can be used to predict customer interest in ads. Consumers are beginning to be wary of cookies, though, and some companies, like Google, are beginning to phase them out of their web browsers.

Third party data aggregators and brokers are increasingly drawing data from multiple sources to provide personalization attributes.[13] They might combine, for example, a consumer's web browsing history with credit card purchases, social media activity, email domain name, type of device used to access the Internet, and other characteristics. The widespread availability of these types of data make it easier to create sophisticated machine learning-based personalization models. Unfortunately, they will probably also hasten the backlash against personalization.

In addition to the attributes of customers and products or services, successful personalization models often include contextual factors about the transaction or offer—including the season, time of day, or specific customer location—that can influence the nature of the offer. Companies may also decide to make offers

only on products that are in stock at a local store or through e-commerce purchases. Only traditional supervised or deep learning models allow for very granular segments using multiple contextual and offer-oriented features to make predictions, so many factors, indeed, that no human could keep them all straight.

In addition to the lack of adequate data, another traditional constraint on marketers seeking to personalize for greater customer engagement is now being eased. AI companies have begun to automate the creation of machine learning models (called automated machine learning, or AutoML) and ongoing maintenance (machine learning operations, or MLOps), making the benefits of machine learning more accessible to non-data scientists. Business or marketing analysts who have some quantitative orientation and who understand customers and markets are, in many cases, now able to create personalization and other types of models using machine learning. They can also ensure that the models do not drift, but continue to effectively predict customer behavior over time using MLOps systems. At Kroger and 84.51°, for example, insights specialists work alongside professional data scientists to create machine learning models, and their greater business insights sometimes make their models more useful than those of data scientists.[14]

Personalization Model Types
Companies can use many machine learning models for personalization, see Figure 1. The most sophisticated AI companies typically combine multiple types for different circumstances. The most common type is supervised learning models, in which systems are trained on labeled data and then used to make predictions.

One example of a labeled outcome is whether or not the customer bought a particular type of product. The model uses that to make a prediction of how likely the customer is to buy it in the future and what factors are statistically associated with that buying behavior. Companies have used traditional supervised ML models, without deep layers, for many years in personalization. Among the companies using supervised learning, among other methods, for personalization are Disney Parks and Resorts (using Magic Band data for personalized itineraries), Netflix, Nike, and Instagram.

Some companies have begun to employ multi-layer or deep learning neural network models for personalization.[15] Among the early adopters of this approach was Dynamic Yield, which was acquired by McDonald's to support its personalization efforts (and those of other companies as well) in 2019, and then sold to Mastercard in 2021. These models require more data and are less readily interpret-ed than traditional ML models, but can often supply more accurate predictions. In addition to McDonald's, companies using the Dynamic Yield personalization technology include Lands End, PacSun, Sephora, and Forever 21. Netflix has developed some of its own approaches and models using deep learning for personalization.

Unsupervised machine learning models can also be used to identify segments or clusters of like customers on several features or variables, and can thus also support personalization.[16] Segments, however, are not a highly precise form of personalization. These models also require extensive data and don't necessarily provide higher value than traditional analytical approaches to data-based segmentation, such as factor analysis or K-means clustering. There are many discussions of unsupervised learning for customer segmentation in the academic literature, but these methods are not widely used in business personalization.

In order to employ multiple models and choose between them, some leading companies are using reinforcement learning, sometimes in combination with deep learning, to find the models best at optimizing longer-term rewards such as a set of clicks or conversions over time.[17] AI systems can automatically evaluate and compare different reinforcement learning models. Netflix, for example, uses reinforcement learning to optimize customers' long-term satisfaction with its entertainment content.

Finally, many companies combine experimental results with other models. Using A/B and multivariate testing approaches helps them to understand customer preferences, particularly for personalizing online content. The outcomes of such experiments are causal rather than correlational, which often provides marketers with a higher degree of certainty about outcomes. Stitch Fix, for example, has built a centralized experimentation platform which uses a variety of experimental

Figure 1
Types of AI Models Used in Personalization

Types of Model	Definition	Strengths	Weaknesses
Supervised learning	Training models with labeled outcomes	Prediction	Labeled data requirement
Unsupervised learning	Requires little or no model training	Clustering like customers	Segmentation only; not better than other methods
Deep learning	Supervised multi-layer neural network models	Adds precision to some predictions	High volume of data required; low interpretability
Reinforcement learning	Maximizes longer-term objective with rewards	Prediction over time	Requires data on sequence of actions
A/B and multivariate testing	Compares two or more online alternatives	Causal outcomes	Becomes complex with multiple arms and when combined w/ML

designs to optimize the effectiveness of clothing recommendations.[18]

An Example of Successful Hyper-Personalization

One highly successful hyper-personalization is Starbucks' AI platform, Deep Brew. Starbucks had historically relied on baristas in physical stores to build customer engagement, but this became increasingly difficult as customers began to order in advance through the company's smartphone app, and was impossible when stores were closed during the COVID pandemic except for drive-through orders or drink pickup. The company launched Deep Brew, which was initially focused on English-speaking markets in the US, Canada, and UK, in 2019. It included not only personalization functions for the Starbucks app, but also assisted workers with administrative activities in the stores, such as reorder points for supplies and labor scheduling.

Deep Brew was Starbucks' first major foray into machine learning, and the company formed a small dedicated team of data scientists to prototype the recommendation engine and other models. Starbucks had massive amounts of data on past customer purchase patterns, which it applied to emailed promotions, in-app featured products and discounts, and games and prizes.

The team concluded that deep reinforcement learning would create the most powerful models for recommendations. This method tests alternative models against each other using an A/B testing approach. In this case, it optimized the criteria of total revenue from a sale and the likelihood of the customer buying additional items beyond their normal purchase. Their goal was to personalize across all touchpoints and channels. Starbucks now uses these models to create more than 10 billion hyper-personalized recommen-

dations a year. The models learn rapidly and continuously from new data.

Now data and customer behavior-based recommendations include likely customer preferences like vegetarian food, price sensitivity, tea vs. coffee, baked good preferences, and more. Recommendations also personalize the drive-thru experience, using not the customer identity, but instead contextual factors such as location, time of day, and weather. During the pandemic, when drive-thru lanes were the only way customers could get Starbucks products from stores, data scientists added a feature assessing the length of the drive-thru line. When lines were uncomfortably long, the drive-thru screen recommended drinks that were easy and quick to prepare.

Within Starbucks, Deep Brew is viewed as successful not only for personalization but for its other administrative functions, and as a driver of engagement with customers who increasingly order through the mobile app; about a quarter of all orders now go through that channel. Kevin Johnson, Starbucks' CEO in 2021, credited Deep Brew with boosting Starbucks' same-store sales, mobile app sales, and drive-through sales during the COVID pandemic period.[19]

The Deep Brew AI project wasn't easy to accomplish, though. Some managers, used to an intuitive and consensus-based decision culture, were not initially receptive to personalization based on opaque machine learning models. Knowing this, the data science team used agile development and reviewed prototypes frequently, carefully tracking the costs and benefits of adopting each model. They maintained a dashboard of key performance indicators during the development of the new model and process behind Deep Brew. They also focus closely on preventing

a security breach or hack. Senior executives calmed Starbucks front-line associates by assuring them that Deep Brew would not replace human workers, but would instead free them up to develop closer personal ties to customers.[20]

These important management tasks of persuading stakeholders, ensuring security, and demonstrating the method's value were perhaps more difficult than the actual modeling. AI managers at Starbucks also emphasize that the model development process was only a small part of the technology development for Deep Brew. The systems or processes necessary to surround Deep Brew's ML models included configuration, data collection, feature extraction, data verification, machine resource management, analysis tools, process management tools, serving infrastructure, and performance monitoring. Together the company had to put much more time, effort, and investment into these infrastructure activities than into the modeling itself.

While most consumers still seem to appreciate some degree of personalization, they may react negatively to offers that seem to be too personal, based on attributes or activities they view as private.

How Much Personalization is Too Much?

The rapid growth in digital media and tools for personalizing ads and offers has prompted a growing consumer concern about privacy, although it has not yet had a substantial effect on personalization or its related customer engagement in the United States. Some have criticized common approaches used for personalization,

including capturing and analyzing customer online behaviors and purchases, as amounting to surveillance capitalism.[21] While most consumers still seem to appreciate some degree of personalization, they may react negatively to offers that seem to be too personal, based on attributes or activities they view as private. Or they may require a higher level of personalization, more accurately attuned to their needs and desires, in compensation for the amount of privacy they believe they are giving up.

We don't really yet know the limits of personalization. For better or worse, surveillance capitalism is in its early stages, and has until recently been hindered by poor quality data, insufficient data, challenges in establishing a persistent customer identity, and lack of methodological sophistication among marketers. Most consumers probably do not view advertising and marketing offers as attacks on their privacy if they are closely targeted with the products and services they really desire.

In addition, individual consumers will have a different sense of the appropriate balance between privacy and personalization. Some will perceive any loss of privacy as a fair value exchange for more efficient and enjoyable shopping. Others will be placated by companies being transparent about consumer data, and even explanations of why they may have seen a particular personalized ad or offer. Facebook, for example, has offered a Why Am I Seeing This Ad? feature since 2014, and in 2019 added further details about personalization approaches, such as the attributes the advertiser was attempting to appeal to, and the source of the data used for personalization.[22] Still other consumers may want no data-based personalization at all.

Personalization of ads and news on mobile devices has been challenged over the last year by changes in vendor policies. Specifically, Apple announced that beginning in 2022 its device Identifier for Advertisers (IDFA) would require users to opt in. IDFA previously allowed apps to personalize ads, offers, and content based on third party app iPhone activity. Prominent app providers like Facebook and YouTube estimated that this app tracking transparency would cost them billions of dollars in advertising revenue.[23]

The ability to personalize marketing is also governed in part by regulatory constraints. The European General Data Protection Regulation (GDPR) does restrict companies' ability to use data for personalization, although its utility in this regard is limited both because individual consumers don't understand the full implications of their consent for personalization and because enforcement mechanisms are limited.[24] The California Consumer Privacy Act (CCPA), the state-level data privacy legislation that took effect in 2020, allows consumers to opt out of having their data sold to another user. However, CCPA has also had little impact on personalization of advertising thus far, in part because few consumers opt out.[25]

Personalization is continuously evolving in response to increasing levels of digitization and data, advancing methods for artificial intelligence, and the changing perceptions of customers and regulators about the tradeoffs between privacy and personalization. Astute marketers can make increasingly effective use of sophisticated personalization approaches, but they should be aware that the value of and reaction to any particular approach is likely to be limited in time. Because of the difficulties of defining and measuring customer engagement accurately, marketers should also address how personalization affects long-term customer behavior and measure multiple aspects of it. It seems likely that well-executed hyper-personalization, driven by artificial intelligence, can help marketers to preserve customers' time and attention amidst an overwhelming flow of information, increasing customer engagement over time. But this expectation is based on logic and certain limited and narrow empirical findings. We will only know more when both research and practice have been pursued over a longer time and with greater breadth than they have been thus far. ■

Author Bio

Tom H. Davenport is a Distinguished Professor of IT and Management at Babson College, a visiting professor at Oxford Business School, a fellow of MIT's Digital Economy Initiative, and a senior AI advisor to Deloitte. He has published twenty-three books and over 300 articles, consulting for leading companies on data, analytics, and AI strategies, organizational structures, and deployment. He is among the world's top twenty-five consultants, top three business/technology analysts, and 100 most influential people in IT.

Endnotes

1. Weng Marc Lim et al, "Past, Present, and Future of Customer Engagement," *Journal of Business Research* 140, Feb. 2022, https://doi.org/10.1016/j.jbusres.2021.11.014

2. The Net Promoter Score is described in Frederick F. Reichheld, "The One Number You Need to Grow," *Harvard Business Review,* December 2003; the newer research on "earned growth," is described in Frederick F. Reichheld, Darci Darnell, and Maureen Burns, "Net Promoter 3.0," *Harvard Business Review,* Nov.-Dec. 2021.

3. Sven Baehre et al, "The Use of Net Promoter Score (NPS) to Predict Sales Growth: Insights from an Empirical Presentation," *Journal of the Academy of Marketing Science* 50, 67–84 (2022). https://doi.org/10.1007/s11747-021-00790-2

4. Thomas H. Davenport, James Guszcza, and Greg Szwartz, "Using Behavioral Nudges to Treat Diabetes," *Harvard Business Review,* Oct. 10, 2018, https://hbr.org/2018/10/using-behavioral-nudges-to-treat-diabetes

5. Thomas H. Davenport, "The Future of Work Now: Morgan Stanley's Financial Advisors and the Next Best Action System," *Forbes,* May 16, 2020, https://www.forbes.com/sites/tomdavenport/2020/05/16/the-future-of-work-now-morgan-stanleys-financial-advisors-and-the-next-best-offer-system/?sh=7fd496fb7027

6. Nidhi Arora et al, "The Value of Getting Personalization Right—or Wrong—Is Multiplying," McKinsey & Co., 2021, https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/the-value-of-getting-personalization-right-or-wrong-is-multiplying

7. "Corporate Data Responsibility," KPMG, August 2021

8. Thomas H. Davenport and Jeanne G. Harris, *Competing on Analytics* (Boston: Harvard Business Review Press, second edition, 2017), pp. 138-139.

9. James Davey, "Loyalty Pays Off for Tesco as Britons Feel the Pinch," Reuters, Feb. 23, 2022, https://www.reuters.com/business/retail-consumer/loyalty-pays-off-tesco-britons-feel-pinch-2022-02-23/

10. Netflix website, "How Netflix's Recommendation System Works," https://help.netflix.com/en/node/100639, accessed August 4, 2022.

11. "The Netflix Effect: How Netflix's Commitment to Extreme Personalization Is Impacting Customer Expectations," WBR Insights, no date, https://nextgencx.wbresearch.com/blog/netflix-effect-commitment-extreme-personalization-customer-expectations, accessed August 4, 2022.

12. The company is named after the longitude of Cincinnati, where it is based. See Thomas H. Davenport, "84.51 ° Builds a Machine Learning Machine for Kroger," *Forbes*, April 2, 2018, https://www.forbes.com/sites/tomdavenport/2018/04/02/84-51-builds-a-machine-learning-machine-for-kroger/?sh=3bd47a0c64e1

13. For an excellent description of the types of data available from third party providers and how it is used, see Stuart A. Thompson, "These Ads Think They Know You," *The New York Times*, April 30, 2019, https://www.nytimes.com/interactive/2019/04/30/opinion/privacy-targeted-advertising.html

14. Thomas H. Davenport, "The Future of Work Now: AutoML at 84.51° and Kroger," *Forbes*, Oct. 21, 2020, https://www.forbes.com/sites/tomdavenport/2020/10/21/the-future-of-work-now-automl-at-8451and-kroger/?sh=6839782f4b0e

15. Deep learning models are neural networks (themselves a form of machine learning) that use nodes based on features or variables to convert data inputs to outputs such as predictions; they are "deep" with several or many layers of intermediate nodes.

16. Unsupervised learning attempts to find clusters, patterns, or anomalies in a collection of data with unlabeled outcomes.

17. Reinforcement learning is a form of unsupervised learning in which a program attempts to maximize a reward in a particular sequence of situations over time. It can employ a variety of algorithm types. It is most commonly used in games, including autonomous chess, Go, and Atari video games.

18. Aaron Bradley, "Building Our Centralized Experimentation Platform," Stitch Fix website, July 30, 2019, https://multithreaded.stitchfix.com/blog/2019/07/30/building-centralized-experimental-platform/

19. Emma Liem Beckett, "Starbucks' Digital Success Partially Driven by AI Engine, CEO Kevin Johnson Says," RestaurantDive website, April 28, 2021, https://www.restaurantdive.com/news/starbucks-digital-success-partially-driven-by-ai-engine-ceo-kevin-johnson/599182/

20. Jennifer Warnick, "AI for Humanity: How Starbucks Plans to Use Technology to Nurture the Human Spirit," Starbucks website, Jan. 10, 2020, https://stories.starbucks.com/stories/2020/how-starbucks-plans-to-use-technology-to-nurture-the-human-spirit/

21. Shoshana Zuboff, *The Age of Surveillance Capitalism* (Public Affairs, 2020).

22. Iris Hearn, "Facebook Expands "Why Am I Seeing This Ad" Feature to Provide More Detailed Insights to Users," *Impact*, July 15, 2019, https://www.impactplus.com/blog/facebook-expands-why-am-i-seeing-this-ad-feature-to-provide-more-detailed-insights-to-users

23. Matthew Fox, "$315 Billion in Market Value Erased After Apple IDFA Privacy Changes," Business Insider Markets, Feb. 3, 2022, https://markets.businessinsider.com/news/stocks/facebook-meta-stock-apple-idfa-ios-privacy-change-social-media-2022-2

24. Michèle Finck, "The Limits of the GDPR in the Personalisation Context" in U. Kohl, J. Eisler (eds), *Data-Driven Personalisation in Markets, Politics and Law* (Cambridge University Press, 2021) https://ssrn.com/abstract=3830304

25. Kate Kaye, "California's privacy law has had 'no impact' on ad revenues or inventory, but indirect effects could hurt." Digiday website, Feb. 25, 2021, https://digiday.com/media/ccpa-early-impact/

Saving Lives with AI:
Lessons in **Personalization** and **Engagement**

Rex Briggs
MMA Global

Stefanie Friedhoff
Brown University

Erik Lundberg
ArtsAI

©shutterstock.com/vs148

Rex Briggs, Stefanie Friedhoff, and Erik Lundberg describe how the Ad Council rose to the challenge of educating people about the COVID-19 vaccination by using AI personalization to get the right message to the right person. The council's work not only saved lives and reduced hospitalizations, it also demonstrated that AI personalization can increase engagement and improve advertising results.

Since its founding in 1941, the Ad Council has brought together the titans of the advertising industry to collectively do social good. During the COVID-19 pandemic the council embarked on the 'It's Up To You' ad campaign to encourage people to visit 'GetVaccineAnswers.org' for accurate and practical information about vaccinations as well as where and how to get them.

The campaign, part of the council's COVID-19 Vaccine Education Initiative, represents one of the largest public education efforts in U.S. history, with support from more than 300 major brands, media companies, community-based organizations, faith leaders, medical experts, and other trusted messengers. It was crafted to appeal to distinct audiences, and it is arguably one of the most engaging educational journeys a person could take, with detailed information about vaccinations, the effects of mRNA on the body, and the pandemic itself.

It turns out that engagement saves lives. An Ad Council regression analysis of visits to GetVaccineAnswers.org and vaccination rates by Designated Market Area (DMA) found that public engagement with the website increased vaccination rates. For every 300 vaccinations, one COVID-19 death is averted and six hospitalizations are avoided. The Ad Council's campaign motivated about 2 percent of total vaccinations, according to our research, which was separately confirmed by Johns Hopkins University.[1]

This engagement helped to alleviate the immeasurably heavy toll exacted by COVID hospitalization and death, along with their high costs. COVID-19 hospitalization costs $75,000 on average for each patient; the U.S. government puts the value of each life at $11.8 million. Lost productivity due to COVID illness further adds to the cost. In total, the outlay comes to more than $5 trillion per year, or about 20 percent of the GDP.[2]

We wondered if we could improve these results with predictive AI personalization. In theory, it would save more lives, keep more people out of the COVID wards in hospitals, and lessen the economic cost of the pandemic in the U.S. We determined to test our hypothesis by launching an AI experiment.

The AI application

AI personalization brings together two dynamic elements: message and context. The message is delivered within a certain context, whether it is a place, a day, a time, a website, or a device. Both the message and the context will influence the response. For example, an ad featuring a cowboy delivered to people in a rural ranching community might lead to more engagement than an ad showing the Las Vegas skyline. We define engagement simply by the number of conversions, that is people completing

a desired action, in this case seeking more information.

On the GetVaccineAnswers.org website, AI observes different combinations of messages and contexts and their conversion rates. It then learns what underlying factors lead to higher conversion rates to predict other successful combinations. The AI predicts which message in combination with which context will most likely result in engagement. (This approach does not require personally identifiable information (PII) or cookies, which would have been controversial given the political climate.)

This AI-based approach is different from testing specific messages on set groups within the audience because the system can study many more message and context elements simultaneously.

As for developing ads, sometimes termed creatives, the traditional approach is to define segments, consider what motivates each segment, and develop a core message to match. The final output is typically one version of the advertisement. Sometimes there are several versions, or a few variations on the same theme, such as English and Spanish language versions.

In contrast, AI-assisted content creation yields hundreds or thousands of different versions, generated by a combination of humans and AI. The role of the human is to curate the ingredients the AI can use to assemble the ads. The role of the AI is to learn which combinations of message elements work best for different audiences in different contexts, and to automatically apply this understanding to optimize and increase engagement rates. The new way should be more effective at raising engagement than the old way.

The AI process

For our experiment, we worked with the nonprofit Immunize Nevada and its ad agency Estipona, with

ArtsAI as the AI tech partner.[3] Our first effort started with three images, and three taglines — a total of nine combinations. (If we were doing video, we could have selected female and male voice-overs, different music, and more.) Our pro bono creative partners on this project, Colleen Watny, Javad Ahmadi and ArtsAI, started with the following imagery:

1) An empty church
2) The Las Vegas strip skyline at night
3) A horseback rider carrying an American flag at a rodeo

We devised three headlines:

1) Let's get back to this
2) Let's get back to normal
3) Are You A Parent? 45,119 kids under 18 lost their parent to COVID-19

Granted, the first two are not that different. We were curious whether the word 'normal' would help, hurt, or make little difference.

The AI predictive personalization technology sorts out which message combinations produce higher engagement with different groups of people. It optimizes for differences in personal context including the content on the website, the viewer's physical location (zip code, DMA, rural vs. urban), time of day, day of week, whether they are using a mobile or desktop device, and many more factors to determine the experience of a person receiving the message.

It works as follows:

- The purpose of personalization is to use a given person's context to serve them the variation of a message that is predicted to generate the highest engagement. ArtsAI uses a hierarchical approach to assembling the message with AI to match a given person and their context.

  At the top level, there are templates that define the layout of the message and its included elements. The next

level is the included elements, such as the image, headline, and calls to action. After the layout, imagery is generally the most important determinant of customer engagement by context. Headlines and calls to action can also make a difference in engagement rates.

- At the start of the campaign, when the volume of observations for the AI to train on is small, ArtsAI starts with clustering techniques (unsupervised machine learning) to select the creative template that is performing best with a particular cluster. Each impression is placed on a vector in the dataset using one-hot encoding. One-hot encoding, also known as one-of-K scheme, automatically converts the categorical data, such as the type of content someone is experiencing, into Boolean data that predictive AI can readily process.

  ArtsAI also clustered this space using K-modes to reduce the dimensionality so the AI can find underlying patterns more readily. For each cluster, the AI finds the best message template and serves it to the cluster. At this early point in the AI learning process, the message elements are randomized.

- With the accumulation of impression volume, the AI learns which images, headlines, and calls to action best engage different types of people, based on the context data the AI observes. The AI switches to supervised machine learning (ML). Specifically, the supervised ML is logistic regression algorithms because of the small size of the model, shorter training and operational time, and absence of need for calibration. AI generates models

for each creative element, starting with the top-level template and ending with each specific creative variation.

AI assisted content creation yields hundreds or thousands of different versions, generated by a combination of humans and AI.

From this process the AI determines the calibrated conversion probability it expects to get after showing a template and a specific creative variation. The question put to the AI is this: What is the probability that the person will engage if they see this particular template and specific creative variation now, in this context?

The creative, or ad, is built individually for each person by first choosing a template and then selecting nested variants of elements with the highest predicted probability of engagement. If the model for a specific element does not pass the quality threshold, the AI serves this element randomly, thus gradually increasing its knowledge.

In essence, ArtsAI's personalization selects the template and elements with the highest probability of conversion using its prediction for each model. AI personalization is different from digital ad server optimization. It is common to use ML to optimize ad delivery, but the ML used by ad servers does not treat the message itself as a hierarchical association of templates, images, calls to action, and so forth. Rather, the ad server treats each version as a separate entity because it is not assembling the elements itself. The result is much less optimization than is possible when messages are deconstructed

Figure 1: Immunize Nevada vaccination campaign's image and headline combinations

	Image: Rodeo and Flag	**Image:** Las Vegas	**Image:** Empty Church
Headline: Let's get back to this			
Headline: Let's get back to normal			
Headline: Are You A Parent? 45,119 kids under 18 lost their parent to COVID-19			
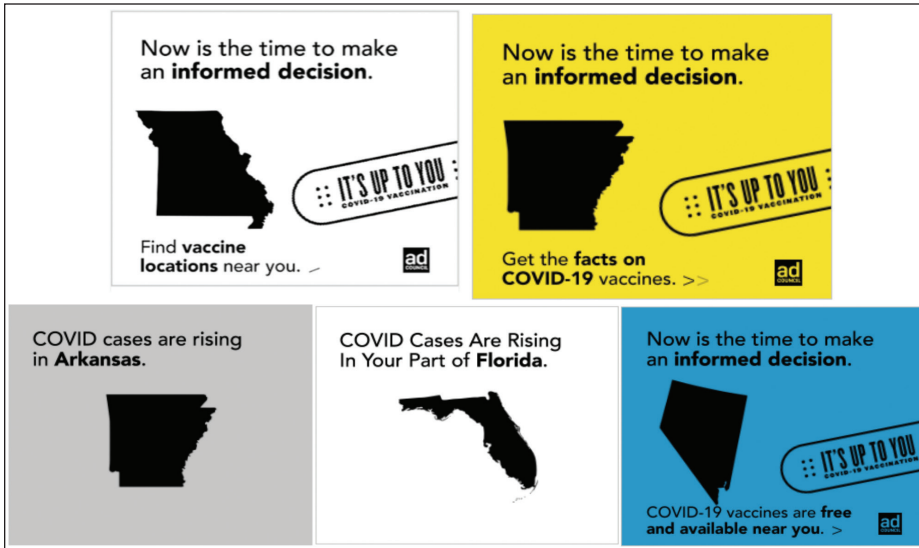
and the AI assembles them in real time.

From nine to more than two hundred versions

After seeing a 43 percent increase in engagement rates in the Nevada experiment, we expanded the effort to the six states with the lowest vaccination rates. Our goal was to raise vaccination participation by giving people vaccine facts to combat misinformation. The 'States' campaign map shows the state where the ad was delivered. The AI then selected from four background colors, three headlines, and three calls to action for a total of 216 versions (6 states × 4 colors × 3 headlines × 3 calls to action = 216).

Note the commonalities and differences by state, headline, call to action, color.

Figure 3: Examples of ad versions

The Ad Council went on to execute two additional AI-powered vaccination information campaigns, one with 158 versions and the next with 155 versions. Below are some of the images, headlines, and calls to action from these campaigns.

Measurement and results

The Ad Council launched four vaccination information campaigns total, with more than 86 million total impressions, meaning the ad was viewed more than 86 million times with an average frequency of 6.6 views apiece. We applied a randomized control group, by effectively turning the AI off. We could then calculate the incremental lift AI provided in conversion to site visit. We found that AI personalization increased lift by an average of 24 percent and motivated 1.1 million incremental vaccinations. Over 21,000 hospitalizations and a median estimate of 3,503 deaths were averted — we conservatively measured the incremental economic impact at over $43 billion.

These vaccination campaigns probably represent a conservative estimate of the typical impact of this AI personalization technology. Across seven total campaigns measured from July 2021 through the end of 2021 (three vaccination campaigns, two in automotive, one in apparel, and one in health care), covering 168 million total impressions, AI personalization roughly doubles conversion rates from 0.59 percent in the control group to 1.23 percent in the exposed cohort, for a 108 percent increase. (See Table 1)

Lessons learned

What makes AI personalization different from the previous generation of advertising is that the versions delivered to consumers are individually selected by AI, whereas the

decisions in earlier methods were based on rules such as delivering one version on the weekend and another on a weekday.

Rules-based models can be labor intensive to program, especially when there are several layers of rules. In addition, rules-based systems assume that the person writing the rules knows which combinations of image, headline, call to action, etc., will work best for each profile. This rule writer has to determine the set of profile variables that define a person and usually does so fairly simplistically because a greater complexity of combinations becomes too cumbersome.

AI revolutionizes the execution of personalized advertising in two ways: first, by letting the AI use a much wider set of profile variables to create unique combinations, and second, by letting the AI learn how to best assemble message elements to get the highest engagement with each profile.

Over 21,000 hospitalizations and a median estimate of 3,503 deaths were averted — the incremental economic impact was conservatively measured at over $43 billion.

Specific lessons learned from the experiment:

Start small and expand after you succeed. This was ArtsAI's advice for our first use of AI predictive personalization: Pick only three images, and try only three headlines. This combination generates nine versions in total.

Table 1: Results of Ad Council States that Used Randomized Control

	Exposure to Site Visit Rate
Randomized Control Group (AI Turned off)	0.5092%
AI Turned On	0.6336%
Percent Lift: **(AI Turned On - Control) / Control)**	**24%**
Lift Range (Lowest to Highest)	16% to 39%

86,111,934 total impressions, 10 percent randomized control group.

After this experience, we found it easier to develop more than 100 versions for our subsequent vaccine communication efforts.

Make sure message elements, especially images, are different. There is no point in testing, say, several shades of green. Also, make the message appropriate for the device. The message, 'Are You A Parent? 45,119 kids under 18 lost their parent to COVID-19' is fairly complex and generated far more engagement on desktop devices. A simpler message is necessary for engagement on a mobile device. Keep this in mind when developing message copy. The AI will sort out which message to deliver to a person on mobile vs. desktop for engagement, but this approach to AI can only optimize what human creative partners feed it.

Have a plan for success. If the AI model is successful, how quickly can you scale the approach? How will you train your agencies and creative partners to develop ads in components that can be assembled by the AI?

We probably could have saved more lives and further reduced the burden on hospitals if we had been ready to scale AI personalization further and faster, but we did not go in with a plan to scale. Instead, we were performing an ad hoc experiment. AI personalization requires a different approach to creative development, and it was a research team composing the ads. We did not have the lead agency creative teams trained and ready to create more campaigns with dynamic advertising.

AI personalization roughly doubles conversion rates from 0.59 percent in the control group to 1.23 percent in the exposed cohort, for a 108 percent increase.

Ad impressions in the 'It's Up To You' campaign generated by AI were less than 3 percent of all digital impressions served. Overall, the Ad Council's campaign engaged people, saved lives, and reduced hospitalizations. So what if at least 60 percent of the messages delivered had been powered by AI? Projecting the lift from 3,503 lives saved from 3 percent, we estimate that we might have averted as many as 70,000 deaths in the U.S. We are proud of the positive impact we made with AI personalization, but we are left to wonder what we could have achieved by adopting AI personalization earlier and more broadly. ◼

Author Bios

Rex Briggs was the Ad Council's data science consultant for COVID-19. He is a pioneer in digital measurement and invented multi-touch attribution and brand lift studies. He was the first to apply neural networks to website personalization. He has written three books: *What Sticks, SIRFs-Up, How Software and Algorithms are Changing Marketing*, and the upcoming, *The AI Conundrum*. He continues his pioneering work as the AI and data analytics subject matter expert for the marketing trade association, MMA Global.

Stefanie Friedhoff is Associate Professor of the Practice of Health Services, Policy and Practice at Brown University. Friedhoff studies the relationships between information inequities, information needs, misinformation, and health outcomes. Prior to Brown, Friedhoff was Director of Content and Strategy at the Harvard Global Health Institute and led programs and special projects at The Nieman Foundation for Journalism at Harvard.

Erik Lundberg is Chief Revenue Officer at ArtsAI, an AI personalization and digital attribution measurement company. Erik is a twenty-five year veteran of digital advertising and a pioneer in dynamic creative optimization (DCO), which has evolved into AI creative optimization (AICO). Erik can be reached at erik@artsai.com.

Endnotes

1. Briggs, and Brown University School of Public health analysis in 2021, www.globabepidemics.org
2. Becker Hospital Review, October 20, 2021, https://www.beckershospitalreview.com/finance/average-charge-for-covid-19-hospitalization-by-state.html, Statistical Life Value (SLV), USTSA value of $2.3million. U.S. Government guidance for 2021 is $11.8 https://www.transportation.gov/office-policy/transportation-policy/revised-departmental-guidance-on-valuation-of-a-statistical-life-in-economic-analysis, Briggs & Brown University School of Public Health Analysis, 2021.
3. More information about these organizations may be found at: https://www.immunizenevada.org/, https://www.estiponagroup.com/, and https://artsai.com/

After a Stroke, AI Helped Me Learn to Write Again

Mukul Pandya

The Wharton School, University of Pennsylvania

Every year, fifteen million people around the globe suffer strokes. Mukul Pandya, a lifelong writer and editor, describes how recent developments in artificial intelligence helped him to recover his abilities and sense of self after a debilitating stroke changed his life overnight.

Globally, 15 million people suffer a stroke every year. If they survive, stroke victims often find themselves in a dark place once they realize their physical impairments may be long lasting. The mental anguish that usually follows can be just as debilitating. That was my state of being after suffering from a stroke in September 2021. I had just retired as the founding editor of Knowledge@Wharton, the Wharton School's management journal, where I was used to being in the midst of some bustling activity around campus.

Overnight, the stroke caused me to lose the use of the left side of my body. I could no longer walk so I had to use a wheelchair. I began to slur my words. Since my hand sometimes threw food around while I was eating, I wore a bib around my neck at mealtimes. I was in shock. I could hardly believe how dramatically my life had changed in a matter of hours.

In her 1997 novel, "The God of Small Things," Arundhati Roy, the Booker Prize-winning author, writes about how life can change in a day. No one knows the truth of that statement better than a stroke survivor.

What bothered me most was the hellishness of being helpless.

What bothered me most was the hellishness of being helpless. I became overdependent on my family, friends, and caregivers, who were saintly in their kindness, patience, and support. Despite my best efforts at trying to stay positive, I had miserably dark days. The main reason why life felt so bleak was the conviction that three months after my retirement from the Wharton School, my professional life was over.

After more than forty years as an editor and writer, I could neither write nor edit. If I could not be a writer or editor, who was I? It was a catastrophic crisis of identity, which every stroke survivor goes through.

Tap, don't type

I had mostly regained the ability to walk, talk without slurring, and use my hands, although it was still difficult to use a keyboard to write and edit. Relearning these motor skills was a painfully slow pro-

cess, but made possible by new technology, particularly artificial intelligence, or AI.

My first baby steps took the form of using tools from Google and Apple on my laptop and iPhone; the apps would look at the words I had typed and try to anticipate and suggest the next word. For example, if I wrote "I am in the …" the AI algorithm would ask if the next word ought to be "hospital." If that was correct, all I had to do was to tap that word rather than type it. This process worked well, I discovered, for text messages and short emails.

Often, if I mistyped a word, the algorithm would underline it and suggest the correct spelling. I could construct short messages to keep in touch with family and friends around the world, although there was a sort of sameness that crept into these texts. Still, it meant that I did not have to wait for anyone else to type emails for me. Although my wife Hema and daughter Tara had kindly done this for me in the first few days after my stroke, the fact that I could do it myself gave me a small measure of freedom. Some of my agency returned. A ray of light broke through the darkness.

The technical term for the AI technology that makes this possible is predictive analytics. As my friend and former Wharton colleague Kartik Hosanagar, author of a wonderful book titled, *A Human's Guide to Machine Intelligence: How Algorithms Are Shaping Our Lives and How We Can Stay in Control*, has helped me understand, "AI increases the accuracy and reduces the cost of making data-driven predictions."

After more than forty years as an editor and writer, I could neither write nor edit. If I could not be a writer or editor, who was I?

In my experience, this tool got better with use. As the AI algorithm learned which words I liked to use – based on the frequency with which I used them – its predictions improved. Gradually, I was able to write longer email messages, even though each message took an excruciatingly long time to compose.

Many friends, including Kartik, suggested trying out speech-to-text software programs, and I did. As the name implies, these are programs to which you can dictate your messages, and the software turns them into text. This AI technology, called automatic speech recognition or ASR, is similar to the one that many companies use to allow you to speak your preferences during a phone call rather than pressing a number on a touch-tone phone. In addition to recognizing commands given in natural speech, these algorithms are able to convert them into text.

Many people have had positive experiences with this software, but my early efforts ended in disaster. I am not sure if it was because I was slurring my words or if the AI algorithm did not understand my Indian accent, but each text message I got was riddled with errors. This was frustrating; it took longer to retype an error-filled message than it did to compose it slowly but correctly letter-by-letter in the first place. There appeared to be no efficient way to write long-form text. I felt doomed to do double work if I were to depend on this technology. I got tired of it quickly and gave up.

Then I found an intermediate solution. WhatsApp, which Facebook (now Meta) acquired in 2014 for more than $19 billion, had a recording feature on its mobile app: I could press a button and speak into the app to create a voice message that I could send to family and friends on my contact list. This was effective because I could now send voice messages that were several minutes long. Often these were complex explanations of medical issues I was dealing with, and I did

not have to face the hassle of the AI algorithm miscommunicating what I wanted to say.

WhatsApp's privacy features also meant that I could speak freely about my health. My friend Rohan Murty, founder of Soroco, a U.K.-based startup, says: "Until you said it, I never thought that WhatsApp could help somebody who's gone through a medical condition like this. If I were a product manager, I would have never realized that maybe one day someone will use it like this."

Another advantage was that my communication could be asynchronous. In other words, I could leave messages for friends in different time zones, and they could respond whenever they had the time. This voice technology allowed me to progress beyond short, terse texts and emails, but I still could not write or edit articles.

Just as the frustration was beginning to build again and the darkness threatened to return, unexpectedly, I had a breakthrough. Before my stroke, I had agreed to interview Google's Neil Hoyne about his book, *Converted*, which is about how companies use data to win customers' hearts. I emailed Neil a list of questions I wrote by typing and tapping on my iPhone. He was kind enough to send back his answers as audio messages. I sent those on to my friend Deborah Yao, editor of AI Business, who had them transcribed, and then edited and published the interview.

Someone reading the article in its final form could hardly have imagined how the process had worked. Thanks to kind and compassionate friends, I was able to produce a long article eight months after my stroke. That gave me an immense boost of positive energy. It was therapeutic and helped me keep healing.

The following month, I was able to do a second story about

ransomware and cybersecurity using the same technique, featuring David Lawrence and Kevin Zerrusen, experts from the Risk Assistance Network + Exchange. The glimmer of hope grew brighter.

Deborah, who previously was my colleague at Wharton, told me about the AI software she had used to produce these transcripts. It was made by a company in Los Altos, CA, called Otter.ai. "Have you tried it?" she asked. "It's good." I downloaded it, and that transformed my life.

Speech-to-text on steroids

My use of Otter.ai was initially a bit complicated. Let us say I had to write a 1,500-word article. I would start by hand-writing a short outline of the story, mapping its structure paragraph by paragraph. (If the article was longer, say 3,000 words, I would map out groups of paragraphs.) After that, I used the iPhone's Voice Memos app, which turns the phone into a recorder, to dictate the entire article. I ended up with an audio file that I could upload to the Otter.ai website.

In a few minutes, Otter.ai's algorithms would create an almost perfectly accurate transcript of what I had said and email it to me. I could now copy and paste the transcript into Google Docs, Microsoft Word, or any other word processing program, clean up the text, and have the final draft ready. While I was impressed that the Otter.ai algorithm got most of the text right, what was truly amazing was the speed with which the AI converted the audio file into text. It could turn even a 60-minute interview into an editable transcript in a few minutes.

What made this magic possible? According to my friend Apoorv Saxena, who once worked for Google and now works for Silver Lake, a private equity firm, it was advances in automatic

speech recognition. An influential paper in 2016 titled, "WaveNet: A Generative Model for Raw Audio," radically redefined the way that algorithms turn speech into text. "We have seen next generation speech-to-text being produced in the last three to four years," he says. This is what makes companies such as Otter.ai as effective as they are, thanks to deep learning.

These days, I use a somewhat different process. Otter.ai lets me create my own digital assistant who 'attends' my Zoom or Google Meet meetings. I introduce 'her,' my AI assistant, as a participant in the meeting to my interviewees, asking if they mind if she joins the meeting to take notes. A few minutes after the meeting ends, 'she' emails me a transcript.

I have taken to practicing typing for an hour every day, so that I can edit the transcribed text. It is important to me to use – but not overuse – the AI technology. If I were to use AI to do everything, I would have no incentive to keep working at strengthening my hand and the neural connections between my brain and fingers. It would simply transfer my dependence from humans to digital technology.

While the AI algorithm that Otter.ai has developed is impressive, it isn't perfect. It gets many things right, but occasionally it gets things spectacularly and hilariously wrong.

While the AI algorithm that Otter.ai has developed is impressive, it is not perfect. It gets many things right, but occasionally it gets things spectacularly and hilariously wrong. For example,

I was recently working on a document in which I had to quote my former Wharton colleague Raghu Iyengar. Otter's transcript turned his last name from Iyengar to "anger" and his first name from Raghu to "Rachel" – getting the name, gender, and nationality wrong. So it still has some ways to go. Fundamentally, though, it has given me a tool with which to resume my writing and editing, and in many ways, to reclaim my identity.

Human-AI collaboration

As I think about the process that has made this transformation possible, I realize that it has to do with structuring human and AI collaboration the right way.

The work begins with a human process (I think of the interview topic, select the right expert, and come up with the questions to ask). Next, I turn over to AI the relatively narrow task of capturing the conversation in audio format and turning it into text. It does this at a speed that is unimaginable for even the world's fastest human transcribers. Finally, I take the task back from the AI algorithm so I can edit and eliminate the laughable "Rachel Anger" kinds of errors and complete the work using human expertise.

I focus on doing what I can do better than the AI and leave it to the AI to do what it does best. This human-AI-human workflow process has allowed me to rebuild my professional life.[1]

© Copyright Mukul Pandya, 2022

Author Bio

Mukul Pandya is the founding editor in chief and executive director of Knowledge@Wharton, the online research and business journal of the Wharton School. After retiring from K@W, Mr. Pandya was a senior fellow with Wharton Customer Analytics and AI for Business. A four-time award winner for investigative journalism, Mr. Pandya has published articles in *The New York Times, The Wall Street Journal, the Economist, Time, The Philadelphia Inquirer,* and more. He has written or coauthored four books.

Endnote

1. To read the entire original essay from which this excerpt is taken, go to: https://www.linkedin.com/feed/update/urn:li:activity:6991399104807346176/

©shutterstock.com/3rdtimeluckystudio

AI + Human Is the Essential Formula for Customer Engagement

Chetan Dube, Amelia

As AI systems for customer engagement become ever more prevalent, it is vital that business leaders understand how to design and use them. Chetan Dube explains why the AI + human formula is essential to high-quality, long-term customer engagement.

Consumers rarely grasp the extent to which businesses (as well as government agencies and a variety of organizations) are rapidly embedding artificial intelligence (AI) systems, in particular to increase customer engagement. An Accenture survey of the world's largest companies found that 75 percent of companies had integrated AI into their strategies.[1] This finding is a comfort since I've been arguing for decades that future success in business — not to mention the prosperity of our planet — depends on building a workforce enabled by AI to address constantly evolving macroeconomic, societal, and labor issues.

Yet in their urgency to bring AI to customer engagement, companies often miss an important consideration. They rush to take repeatable customer service tasks out of human hands, pass them on to AI (usually in the form of interactive digital agents), unveil the technology to users, and then sit back and wait for what they expect will be positive results. While those are good intentions, and such an approach can be initially successful, it's easy

to overlook how essential a digital + human formula is to high-quality, long-term customer engagement. Utilizing either humans or AI without the other will ultimately prove ineffective.

AI systems and humans must work together as a cohesive team to create personalized, ongoing user engagement that can engender repeat business, brand loyalty, and high customer satisfaction over time. Expecting AI to go it alone on the road to better customer engagement will lead companies to disappointment. Some have already discovered this and gone on to successfully pair AI and human efforts.

The AI + Human Equation

Many companies' efforts to combine humans and machines to refine customer engagement are in flux. If a future workforce is to thrive, humans should not be primarily involved in repeatable work that AI and automation can do more precisely.[2] Likewise, machines can provide data analytics and input into a process, but humans should make the decisions about personalized customer care. Both precepts are essential, but many organizations invest in automation without duly considering the critical role that people, with all of their exceptionally human skills, must play. Essentially, they set it and forget it, relying solely on AI systems to make a real impact on their business.

There is a better approach. It hinges on interdependence between humans and AI, with each working to create the best customer experiences possible, blending their skills, processes, and components. Although each organization uses AI uniquely, certain common policies will help them to build an optimal AI + human balance.

Human Competencies and Skills

For the formula to be effective, a company's workforce must fully understand its Information Technology (IT) landscape. The IT teams, who are responsible for applying IT systems to securely complete specific tasks, such as transfers between bank accounts, must understand the business processes and workflows so they can ensure success from start to finish. Workers need to learn how AI can create and accelerate processes and outcomes for themselves and external audiences, such as customers or partners, so they can apply those capabilities appropriately.

Most importantly, humans should realize that when they deploy an AI system, they do not abdicate their own responsibility for achieving positive results. AI can independently perform many functions incredibly well, but it does not do so alone. Human workers must monitor, analyze, and study the AI's performance and outputs, constantly identifying ways to improve it or assign it new jobs, often according to the AI system's own recommendations.

Finally, for humans to partner successfully with AI, they must relinquish their fear that AI will completely take over their responsibilities, or "steal their jobs." Yes, AI may take on tasks that were previously done by humans, such as resetting passwords, tracking orders, or scheduling appointments. But then, who among us would want our entire careers to consist of performing such tasks, when far more challenging and rewarding jobs are emerging every day?

Placing repetitive tasks in AI's hands allows people the time to learn new skills and responsibilities, delivering higher value to their organizations. This newfound freedom empowers people to use their uniquely human abilities such as creativity, insight, inspiration, empathy, and so forth. These abilities are out of the reach of AI systems and are essential to developing innovative products and services.

Put another way, humans should eagerly embrace the new work possibilities that AI presents.

AI Competencies and Skills

To maximize the benefit of an AI + human formula, the best route is to use AI that is designed to closely approximate human qualities, but is bolstered by machine speed and the ability to scale for high volumes on demand. We have taken this approach at my company in developing Amelia, a conversational AI tool that automates back-end systems and processes, communicating in the preferred language of users through a human-like digital agent over any channel, including voice, text, mobile apps, and platforms such as Microsoft Teams. (I'll pause here to proudly note that Amelia is regularly recognized by third-party analysts as the industry-leading digital agent.)[3]

An effective AI system requires a combination of algorithms, machine learning, analytics, workflows, business rules, integrations, and some method of self-learning, in which the system makes improvements and acquires new skills over time as it completes tasks, in much the same way that a human does. In fact, we've taken this approach to modeling Amelia's "brain," designing her mind to operate very similarly to a human's. For example, we have provided her with:

- **Natural Language Understanding** through a combination of deep neural networks and natural language data sources so she can contextually understand and interpret simple and complex multi-sentence requests.
- **Semantic Memory** to store facts, concepts, and the associations between them, and apply them to conversations.
- **Episodic Memory** to recall previous conversations and apply that knowledge to new interactions.

Figure 1: An Organizational Framework for the AI + Human Formula
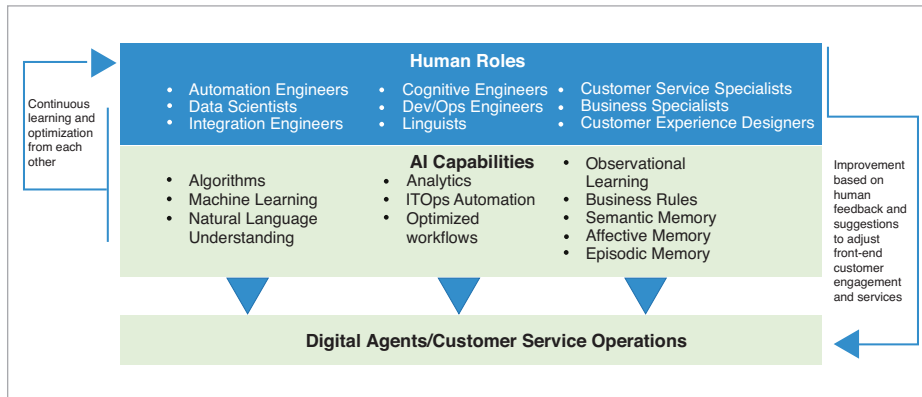
Source: Amelia

- **Affective Memory and Sentiment Analysis** techniques which allow her to model users' emotions, mood, and personality. Essentially, she can learn how to compose responses that are emotionally resonant within a specific conversation, such as offering comfort to a person who was just involved in a car accident.

> How these human and AI capabilities are blended depends on a given company's ability to invest not just in technology, but in people with specific skills, so it can put an optimal organizational structure in place.

How these human and AI capabilities are blended depends on a given company's ability to invest not just in technology, but in people with specific skills, so it can put an optimal organizational structure in place. Although there are variations between companies, the figure above shows an example of a high-level framework. Note how the human roles are not strictly in IT; some are business and customer service specialists who can use their expertise to influence customer experience design. Additionally, human customer service agents can work in tandem with digital ones; when a digital agent is unable to complete a customer service task, it can hand off the issue, along with information and context to avoid loss of time and data, to a human agent for personalized support. What's more, robust digital feedback channels allow practically any human involved, employee, customer, human agent, partner, or even casual website visitor or app user, to contribute to constantly improving customer engagement.

Results from the AI + Human Approach

Results from a small but growing number of companies are beginning to demonstrate the power of the AI + human approach in customer engagement. Companies deploy digital agents, armed with conversational AI and automation to the front lines of customer service. In analyzing the experience of 40 clients that have adopted an AI + human approach with Amelia's digital agent and automation solutions, using actual and estimated metrics and results, I found that several points stand out:

- *Containment:* An AI + human approach is performing well in terms of call containment, allowing customers to resolve their issue without escalating to human agents, whose time can then be reserved for the most critical calls.
- *Intent recognition:* Digital agents (powered by Amelia) have excellent rates of recognizing intent (average 90 percent), accurately determining what a user/caller wants to achieve by contacting a company by voice, text, or mobile app. This ability contributes to successful containment (average 62 percent), and to high rates of call resolution (average 75 percent).
- *Customer Engagement:* An AI + human approach allows companies to bring online new means of engaging with customers that were previously impossible or unfeasible.

The forty clients I analyzed represent a diversity of leading global and regional businesses that generate millions or billions in revenue in various industries including banking, insurance, financial services, technology, and retail. Each required a system that could accurately handle thousands of calls and inquiries each month with excellent customer service.

Containment: Many companies use Interactive Voice Response (IVR) systems for menu-based phone assistance: "Press 1 to speak to a representative." Users are "contained" within the system to complete interactions without the need for human assistance which can be time-consuming for callers, asked to hold for a human agent, and impractical for companies that may have too few representatives. As anyone who has attempted to navigate these systems can attest, many IVR systems lack the intelligence and information to handle anything more complex than

basic FAQs or simple bidirectional transactions (one input in, one output out).

Many of our clients augment their IVR systems with our digital agent platform, armed with natural language and context switching to handle more than one issue at once, with an eye to raising containment rates. One of the immediate benefits of deploying a digital agent platform like Amelia is that customers' hold times can effectively be cut down to zero by the inherent availability and scalability of the technology. In addition, when looking at actual and estimated results from our forty clients, I found that many have excellent containment rates, with an average of 62 percent, a median of 73 percent, and a range from 12 percent to 100 percent. This result does not mean that digital agents are handling all calls, only that they are playing their assigned role of keeping customers contained within automated systems. Amelia can complete entire interactions so as to reduce call volumes and reduce or eliminate IVR wait times. When a digital agent is unable to complete a task, Amelia seamlessly transfers calls to human agents, ensuring a smooth handoff without any lost data or context, so a human can pick up right where Amelia left off.

This kind of impact on containment can produce ancillary benefits in overall customer satisfaction. One telecommunications provider uses Amelia as a live digital customer service agent to handle 100 percent of its mobile call volume.[4] A positive Net Promoter Score (NPS) is a critical measure of a provider's investment value to internal stakeholders. This provider's digital agent recorded a NPS that was, on average, 16 points higher than that of an approach that used only human agents (see figure 2 titled Mobile), a sizable change in customer sen-
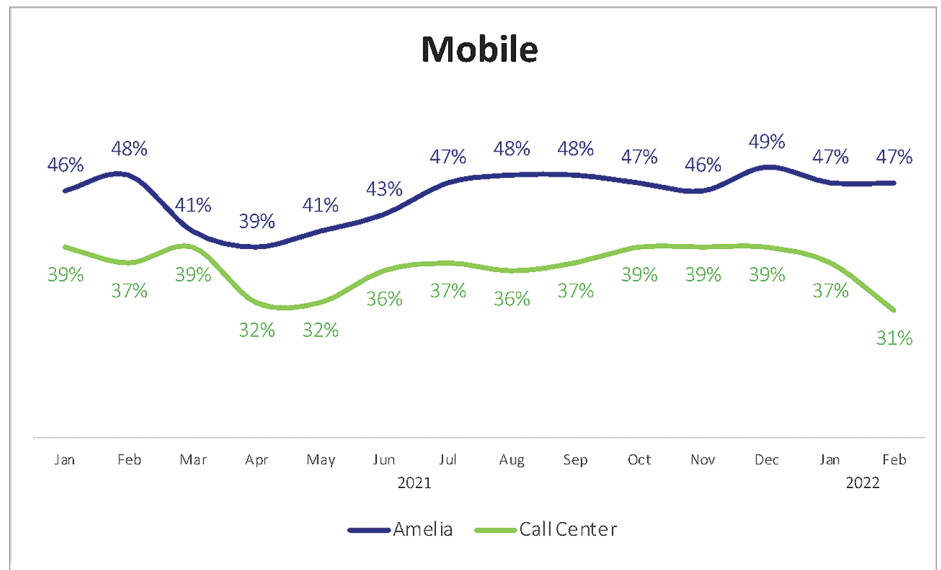
Source: Amelia

timent. This data also indicates that human agents, once they are no longer burdened with the most routine calls, can be retrained and tasked with higher-value functions, thus raising their NPS score.

Intent recognition: A digital agent is a reflection of how it is trained, so a successful AI + human approach starts with AI's ability to clearly understand a user's intent and provide exactly what they want or need. This requires, among other things, an ability to understand not just languages but how people actually communicate, whether it be in full sentences, colloquialisms, or short phrases, words, and sounds.

A digital agent which cannot recognize intent is incapable of resolution, and its interactions with users can go nowhere.

A digital agent which cannot recognize intent is incapable of resolution, and its interactions with users can go nowhere. This is why

the 90 percent average rate of intent recognition I discovered across our deployments to forty companies is so encouraging. And it has a corresponding impact on resolution rates, measuring the ability of digital agents to successfully complete an interaction. Our customers currently record an average 75 percent resolution rate, with a 81 percent median on a range from 50 percent to 93 percent.

Ability to engage new customers: Our clients have based their AI + human approach on specific business challenges, but there is some commonality in the new methods they used to reach current and new customers. AI solutions can be deployed across channels to match user preferences, so many businesses expand customer service beyond phone-based support to include mobile apps, text, and communications platforms such as Microsoft Teams. This multi-channel ability provides true 24/7 service, because digital agents can handle even complex transactions outside of normal business hours.

Many companies also take advantage of AI's multilingual ability (Amelia speaks 100 languages) so

as not to force users to select one language over another, as may have been previous practice because of technology limitations. A small number of companies, like Resorts World Las Vegas, create their own AI avatar or virtual presence for their digital agent, one that embodies their corporate identity and brand, and can demonstrate their commitment to engaging with users.[5] At the resort, the customer experience is dramatically enhanced. Amelia, renamed Red, handles the company's call center solo. At any time, guests can speak directly with Red for assistance with anything including making dinner reservations, purchasing entertainment tickets, ordering room service, setting wake-up calls, and much more.

Or consider Visionworks, an Amelia client and one of the largest eye care providers in the US.[6] Amelia's responsibilities there consist mostly of laborious tasks that were previously done by in-store employees such as scheduling appointments and fittings, providing order status and pick-up information, describing safety and COVID-19 protocols to in-store visitors, confirming insurance coverage, and providing location services.

Another example of expanded customer engagement is Amelia's role at ASICS, the global sports company, where she provides 24/7 support through email, website chat, and social media. Amelia assists ASICS' customers with order and return status updates, dispute case status, return instructions, and shipping information.

The AI + Human Approach at Work

In a typical execution of an AI + human approach, a company harnesses all of AI's intrinsic benefits and then makes subsequent alterations to human jobs, with a focus on employees providing new products and high-value customer service. AI, meanwhile, takes on lower-value roles or critical tasks that complement those of human partners. Digital and human agents mentor one another in coordinated customer service, working toward the same objective. The performance of digital agents spurs recommendations while humans alter user services to improve customer engagement in the long term.

One striking success story of the AI + human formula in action comes from an industry you may not expect: legal services. Our client Kenneth S. Nugent PC, Attorneys at Law, one of the US' largest personal injury practices, uses Amelia to offer AI-enabled, personalized client service by phone, website, or secure client portal.[7] The firm also uses a digital agent to address high call volumes and to free its staff and attorneys from rote tasks so they can focus on bringing positive legal outcomes to clients.

The thirty-year-old Georgia-based firm, which employs more than 200 staffers and forty-seven attorneys, fields thousands of calls and inquiries each month. Before AI the road from initial contact to signed contract was far from seamless and sometimes resulted in clients signing with competing firms. The firm's clients also expressed an increasing interest in conduct business through digital channels, especially after office hours, though the firm had always leaned toward person-to-person interactions.

Beginning in January 2020, the firm piloted our Amelia digital agent.[8] Its primary purpose was to use Amelia as a legal assistant, taking on administrative duties for the firm's intake specialists, receptionists, and legal teams. The law firm programmed and trained its digital agent using a roster of skills, dialogues, and procedures specific to filing personal injury claims. These lessons included basic terminology, the ability to request and record sa-lient information about an accident or injury, and natural language processing to understand user interactions in context, through sentences of varying complexity. Some of this was accomplished by uploading data on claims processing to Amelia, while other, more custom workflows were developed.

Amelia is now deployed on a 24/7 client portal where she communicates via voice or web chat. She is integrated with the firm's practice management and customer relationship management (CRM) systems for creating and tracking cases and for contract processing. Custom-built intake tools also allow her to securely record accident and injury details and upload documents and photos.

In addition to the portal for existing clients, Amelia interacts with new clients via web chat on the firm's homepage, and is the first point-of-contact by phone, collecting initial information before transferring callers to the appropriate case manager. In some instances, a case manager isn't even necessary; for new clients, Amelia can take down initial information, answer questions, and independently initiate a signed client contract.

> The firm's founder reports that staff members are very positive about the technology; they do not view Amelia as subsuming their jobs, only improving them.

Case information is stored in clients' digital files, and the digital agent (which introduces itself as a digital receptionist) is programmed to ask for updated information when a client returns to the portal. Clients continue to communicate

with the digital agent after their initial contact. The firm's founder reports that staff members are very positive about the technology; they do not view Amelia as subsuming their jobs, only improving them.

After the January 2020 launch, Amelia was fully up and running by May. On average, Amelia handles some 15,000-20,000 operator calls, 2,000-4,000 website chats, and 2,500 client portal visits each month. Amelia processes 65 percent of phone calls independently, without referring them to a human operator. Perhaps most impressive, in two years the digital agent has produced almost 3,000 signed client contracts, without any initial human intervention (human case managers continue to perform contract work). Before Amelia was deployed, clients would often contact their attorney or case manager to speak at length about case-related issues such as health records, doctor's appointments, and so forth. Since Amelia is now responsible for those conversations, the firm's teams can focus on finding the best way to litigate a case, and clients still feel valued.

Amelia has virtually eliminated the firm's concerns about losing potential business due to unanswered calls. Before Amelia, about 5 percent of calls – or some 1,000 a month – were not answered. Most firms have a 10 to 20 percent unanswered call rate. Now the firm reports that no calls go unanswered, and Amelia has returned almost 200 hours a month to human operators.

Clients, it turns out, enjoy having a digital agent that is always ready to discuss their case.

Amelia's ability to record and answer questions via web chat also provides users with an effective channel which lies within their comfort zone. Before Amelia, an average web visitor would spend about 1.5 minutes on the portal; that number has increased to 7.5 minutes for both first-time and returning visitors, which the firm attributes to Amelia's 24/7 availability. Clients, it turns out, enjoy having a digital agent that is always ready to discuss their case.

The firm initially used the Amelia avatar for its website chat and client portal. Now, Nugent has gone one step further, working with Amelia to create a "digital twin" in his own image. Clients who interact with the firm through web chat or mobile see his likeness on their devices. Nugent has thus become the virtual personification of his own business

In some ways, this approach to AI + human engagement parallels the currently limited use of AI in higher education.[9] To achieve a true AI + human approach we must expand these roles considerably. Just as AI is upending how many businesses view customer-centric care, it can also drive higher education to move from current models of faculty-centered education to learner-centered ones. Interactions between teachers and students can be supplemented with 24/7 digital/human teams. These would pair virtual agents which can converse and interact with students using machine-speed access to curricula, syllabi, research, and other information, with human mentors to provide critical context and help students learn.

Using AI for customer engagement, however, does require designers in different industries to address different considerations. A digital agent's interactions with a banking client or an insurance policyholder would be markedly different than its interactions with a college student or a patient scheduling a medical appointment.

A banking customer would expect exchanges to be prompt and business-like, mostly transactional, and with a degree of empathy appropriate to the issue at hand (e.g. a lost credit card). A person trying to schedule a surgical appointment with a specialist, on the other hand, might be anxious and nervous. That engagement should therefore begin with empathy and comfort before the transaction of setting the appointment occurs. Companies must also be sure to respect various ethical considerations such as the sharing of financial information, health records, or student information.

Preparing for the AI + Human Future

A holistic approach to customer engagement will fuel a company's ability to survive financially and competitively. Businesses that adopt an AI + human method, one that places equal weight to both sides of the equation, will thrive; those that continue to think in more limited terms will see their revenues decrease, their margins shrink, and the confidence of their customers erode.

Beyond customer engagement, the potential benefits of an AI + human approach will be evolutionary, not revolutionary. As its influence grows over time, business leaders should set their short- and long-term expectations accordingly. For example, we expect that our base of some 200 direct clients will reap cost efficiencies of 5 to 30 percent in the first year and additional savings in subsequent years as they continue to optimize their business processes. These savings depend in part on whether these companies expand their use of AI and automation. As rote process work is handed off to digital agents, human employees can be reassigned and retrained for higher-value tasks, creating a workforce engaged in challenging

and satisfying pursuits, and thus to high employee satisfaction and retention.

One of the best indicators of beneficial customer engagement with a digital agent is volume, the number of conversations the digital agent conducts in all channels. Strong performance in these sessions offers a window into a digital agent's efficacy, whether it understands what users are asking, and adaptability, whether it can deal with a wide range of customer inquiries and learn more over time. It also demonstrates users' overall acceptance of the agent, their satisfaction with its abilities and their willingness to engage with the digital agent again.

Looking across thirty-nine clients, I found that the average increase in sessions with Amelia during a recent six-month period was 21 percent and the median 52 percent. This metric was strongly influenced by how long the client had been using the system and thus how well established it was. Some had installed the platform right at the start of the timeframe. Others were longer-term clients reporting month by month increases in sessions. One insurance company reported that, since going live six months earlier, conversations with Amelia

had increased from less than 200 per month to more than 50,000 per month. A financial services company that had deployed Amelia two years previous was still seeing substantial gains in user conversations with Amelia, from 190,000 to 260,000 per month, while maintaining a 90 percent call containment rate.

Implementing the AI + human approach is not risk-free, particularly if a company fails to prepare for how it will profoundly change corporate operations. Leaders who expect AI to be a panacea that will magically counter years of inaction in process or customer service improvement are likely to be disappointed. Similarly, those whose main purpose is to appear innovative, or to react to competitive market forces, are likely to fail in planning AI's proper use or in due diligence more broadly. Bolting AI elements to digital channels or the back end of already inefficient processes is often a pointless exercise. Consider the recent customer service turmoil in the airline industry to see how digital is irrelevant if companies do not address underlying issues.[10]

Despite risks, though, it's imperative that business leaders make the attempt, stepping out of

their comfort zones and having the courage to experiment so they can make this model a reality for their companies. With millions or billions dollars of capital investment at stake, this step is not nearly as straightforward as it may sound. It requires that leaders thoroughly rethink customer engagement, working from a desire to meet customers where and when they are most comfortable.

Finally, companies must properly prepare their workforces for the changes that this model will produce. This preparation requires that leaders involve employees in the planning, designing and deployment of AI + human projects. Employees who participate in the process tend to feel a sense of ownership over the model, recognizing its potential to improve not just the business but their own jobs and livelihoods. They know that they are an essential part of the AI + human approach.

It is therefore incumbent upon companies to trust the technology and their employees in equal measure so that the AI + human formula works harmoniously, producing impactful customer engagement. Any strategy that doesn't account for the profound contributions of both is out of tune and out of touch with where the market is going. ■

Author Bio

Chetan Dube is the CEO and Founder of Amelia (formerly IPsoft), launched in 1998. A mathematician, Chetan was previously an assistant professor at New York University. He is a widely recognized speaker on automation, autonomics, cognitive computing, conversational AI and the future digital workforce. He serves on the boards of numerous IT-related institutions. In 2017, Forbes listed him as one of the Nine Greatest AI Minds. He is also a member of the Forbes Technology Council.

Endnotes

1. https://www.accenture.com/in-en/insights/artificial-intelligence/ai-maturity-and-transformation
2. https://www.mckinsey.com/capabilities/operations/our-insights/the-imperatives-for-automation-success
3. You can learn more about Amelia here: https://amelia.ai/conversational-ai/
4. https://amelia.ai/customer-story/amelia-handles-100-of-mobile-phone-traffic-for-telefonica/
5. https://amelia.ai/customer-story/amelia-enhances-guest-and-employee-experiences-at-resorts-world-las-vegas/
6. https://amelia.ai/press-release/visionworks-hires-amelia-to-enhance-customer-service-and-reduce-costs/
7. https://www.attorneykennugent.com/
8. https://amelia.ai/video/the-amelia-story/
9. https://www.insidehighered.com/digital-learning/blogs/online-trending-now/artificial-intelligence-assist-tutor-teach-and-assess
10. https://www.wsj.com/articles/airline-customer-service-wait-times-11650999776

The Leader's Strategic Mindset:
A Key Factor for AI Success

Pernille Rydén
IT University of Copenhagen

Torsten Ringberg
Copenhagen Business School

Omar A. El Sawy
Marshall School of Business, University of Southern California

Businesses are increasingly using artificial intelligence (AI) to transform their processes and create new ways of engaging with customers. But leaders who project old assumptions about customer engagement onto the breathless use of AI risk creating a dangerous disconnect that will deprive their companies and their customers of value. Pernille Rydén, Torsten Ringberg, and Omar A. El Sawy present a strategic mindset framework to help leaders avoid this pitfall.

AI is currently one of the most in-demand digital technologies for driving businesses forward. Its allure to enterprises lies in its ability to execute tasks that previously could only be done by the human brain, and to do so faster and often better.[1] Because of AI's utility and efficiency, ever more businesses are adopting it. The global AI market is expected to reach $267 billion by 2027.[2] And yet to derive full value from the use of AI, it is important that enterprise leaders consider carefully what mindset should govern their strategic decisions about AI's use, including increasing customer engagement.[3] AI can escape some human cognitive pitfalls but will never be bias-free since its usage is determined by the underlying strategic mindset of business leaders.

Strategic mindset determines how leaders answer the question, 'why AI?' The mindset defines how leaders understand the opportunity of AI and how they apply it to everyday business interactions. But a mindset is a human's black box where habitual mental programming formed by past experiences and assumptions produce a take-it-as-a-given reality.

March and Simon[4] found that leaders who rely on simplified mindsets are hampered by bounded rationality when they face complex problems in the market.[5] Their results highlight not only why leaders are suboptimal decision-makers, but also why their decisions are heavily biased by their underlying mindsets.[6] So leaders' assumptions about AI deployment and future business scenarios are influenced by their existing mindsets, which may hinder their enterprises, efforts to explore the new opportunities of AI.[7]

Finding the right 'why' can be challenging. It depends on what leaders want to achieve, who their intended customers are, and what needs AI can help them meet. Many leaders tend to skip these steps and instead focus immediately on what AI they should use and how.

Four Mindsets

The strategic mindset framework helps leaders to identify which strategic mindset will best suit their use of AI to increase customer engagement by illustrating four ways of engaging with customers.

The left side of the model represents classical modes of marketing interaction. Quadrant A represents the cold sale, the analytical use of AI to push products and services to consumers. Quadrant B reflects observing consumers' preferences, needs, and motivations, using AI to draw new insights to make customer engagement more meaningful. The right side of the model represents modern market thinking, with types of engagement that involve customers more. Leaders in Quadrant C use AI to glean new and visionary insights from consumers by inviting them to collaborate and innovate. And Quadrant D shows a more empowering and sustainable approach to business engagement with customers, using AI to entice consumers to make informed decisions that have a positive effect on their own lives and on their communities and the planet.

Enterprise leaders usually have an over-arching mindset that frames their customer engagement strategy, but aspects of the other three mindsets often support that governing mindset. This structure affects what AI tools the company uses and how. Since the mindsets are not mutually exclusive, their influence on decisions about AI should be aligned to serve a common strategic direction.

Leaders must therefore become consistently aware of how their mindset limits their outlook and, consequently, their ability to effectively use AI in new ways to engage with potential and existing customers.

For instance, the goal of using AI to deliver a great customer experience may be inadvertently overruled by other priorities such as using AI to save money, cut lead times, boost process optimization or data compliance, and build cyber security. Leaders must therefore

Figure 1: The Strategic Mindset Framework[8]
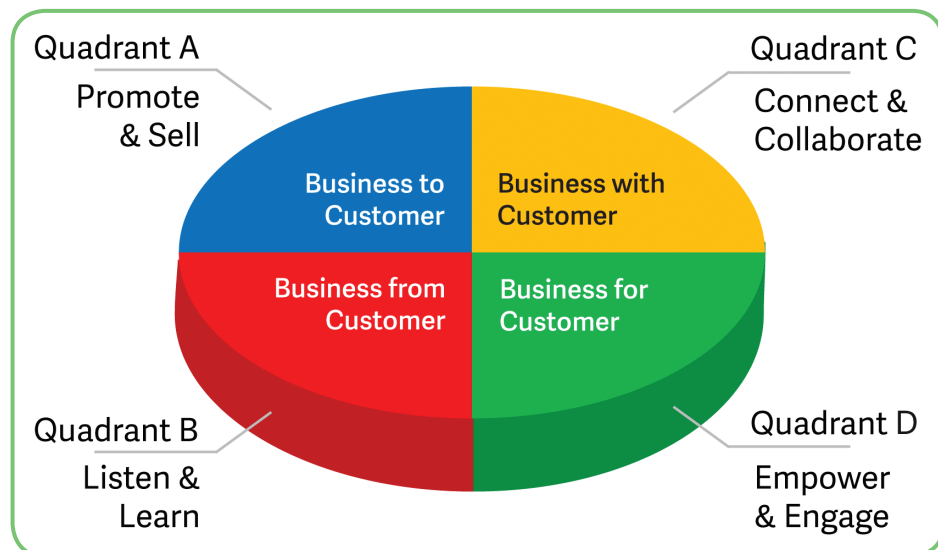
Quadrant A
Promote & Sell
Business to Customer

Quadrant C
Connect & Collaborate
Business with Customer

Quadrant B
Listen & Learn
Business from Customer

Quadrant D
Empower & Engage
Business for Customer

become consistently aware of how their mindset limits their outlook and, consequently, their ability to effectively use AI in new ways to engage with potential and existing customers.

The Promote and Sell Mindset

Leaders who rely on a classic transactional business to customer interaction logic use AI to automate marketing and sales operations, increasing sales, enhancing marketing, improving operational efficiency, to reach more customers faster.[9] AI can replace the early stages of the traditional sales funnel, fortifying the sales pipeline and online stores to boost efficiency, managing customer relationships in real time, identifying opportunities to cross-sell, and predicting customer behavior.[10]

Facebook and Google, for example have multibillion-dollar ad businesses that offer AI capabilities for ad management.[11] AI support chatbots, algorithmic trading,[12] and entertainment recommendations to make ongoing and smarter offers that increase customer purchases.[13] Bots that directly evaluate and sell to customers reduce the need for expensive, commission-earning sales agents.[14] Amazon has benefitted enormously from using AI applications across e-commerce, logistics, and warehousing, ranging from immediate recommendations to Alexa-enabled voice shopping.[15]

Walmart uses AI to help customers through its personal shopper program that suggests the best substitute for an out-of-stock item.[16] Macy's is testing cognitive AI technology powered by IBM's Watson AI to help customers navigate its stores. The smartphone-based assistant Macy's on Call can answer questions about where to find products or brands, services, and facilities in the store.[17] By using AI to replace some human shop assistants, companies can optimize their sales and make their service more reliable.

Enterprises driven by the promote and sell mindset can tap into myriad connected digital devices and platforms. AI can then use real-time behavioral data to understand what customers consume and when, rather than the deeper needs that drive their consumption. The resulting value chain is tightly controlled, which ensures consistency, security, and value, but it does risk misalignment between offers and what consumers truly need.

The Listen and Learn Mindset

To develop value propositions, creating offers that appeal to customers, it is increasingly vital that enterprises know what those customers really want.[21] They can use sophisticated AI to *listen* to customers and *learn* about their needs and behavior in order to make targeted offers. AI can pick up and employ emotional cues

Table 1: Business-*to*-Customer AI Engagement	
REAL-TIME OPERATIONS	**CUSTOMER CONVENIENCE**
· **Optimize production, distribution, and sales:** Identify and anticipate product shortages and peaks due to bottlenecks in production or spikes in demand.[18]	· **Faster and cheaper products and services:** Monitor customer behavior through instantly available live data feeds to improve services and to drive sales in enterprises.[20]
· **Reach more customers faster:** Marketing campaigns using large sets of integrated customer and product data and employing analytical models can target large numbers of customers instantly.	· **Available offers and communication:** Identify the priority of offers, based on factors such as customer lifetime value and offer value, and pre-assigned business rules, to identify the best offer for the customer.
· **Increase product and service reliability:** Deliver promotional and sales offers across a wide range of automated and face-to-face channels, such as the web, mobile, or call center agents.[19]	· **Less friction frees personal resources:** Optimize self-service opportunities and enhance real-time, data-driven digital customer experience through virtual customer assistants and chatbots to drive marketing automation.

and language in a human-like way that brings customers closer to the brand. AI can also analyze how customers feel by instantly identifying emotional cues in the words they use, which may generate insights into their underlying motivations.[22]

For example, Kore.ai is a conversational AI software company. Its omnichannel virtual assistant, BankAssist, uses AI-enabled interactive voice response to communicate with customers in their own personal context. It can also personalize and automate customers' relationships with businesses through highly accurate natural conversations with maximum containment, that is, the fewest possible human interventions.[23] Because the system is integrated with major core banking, cards, digital banking, and bill pay, customers can start conversations in one channel and complete them in a different channel without interruption or loss of context.

Companies can scrutinize the entirety of the customer experience using AI rating tools, which use survey questions specific to each company to rate the interactions for which there is no customer feedback. Enterprises like Home Depot, JP-Morgan Chase, Starbucks, and Nike use AI to monitor customer experiences and behavior and to generate real-time insights which ensure personalized, seamless, omnichannel customer experiences.[24] They can also use this data to intervene quickly when things go wrong, achieving effective service recovery.

Facebook, Instagram, and Tik-Tok tap into personalized social media data streams and use AI to learn how to create relevant offerings that are rooted in both social and emotional needs and behavior. Siri and Google silently listen to users through smartphones and track movements and location. Telecom giant Comcast uses Pointillist, an AI and machine learning customer journey analytics software, to rapidly diagnose problems and identify how to increase customer satisfaction and recover quickly from failure.[25]

So, by ensuring that their internal capabilities and consumer insights are better than those of their competitors, enterprises can increase customer satisfaction, creating unique value propositions. Moreover, customers who receive relevant propositions are more likely to share their personal data and remain loyal.

The Connect and Collaborate Mindset

Customers who prefer to further engage with enterprises will appreciate a connect and collaborate approach, in which enterprises benefit from cocreative and crowd-sourced business-*with*-customer engagement. Here, AI builds knowledge networks as well as personalized and informal interactions, increasing the flow of ideas and knowledge between consumers and enterprises, which allows those enterprises to efficiently create, disseminate, and revise its offerings.

Table 2: Business-*from*-Customer AI Engagement	
REAL-TIME RESPONSES	**CUSTOMER RELEVANCE**
· **Instant personal customer data insights:** Identify, gain, develop, and retain profitable customers, learning what they need by generating, monitoring, and measuring their sentiments and preferences. · **Optimize and segment processes to fit customer preferences:** Content recommendation algorithms personalize the information on social media, search engines, smart speakers, and other applications. · **Integrate immediate customer ratings:** Automate rating processes with AI rating tools for more actionable analytics.	· **Customize recommendations and digital products and services:** Offer highly personalized predictions and recommendations. · **Customize engagement to strengthen customer relationships:** More accurate segmentation and computation of individual preferences and anticipatory interaction. · **Fast feedback from customer input increases relevance:** Respond to customers' increasing demand for faster and better services.

AI allows enterprises to actively solicit customer's input, further improving their offerings and deepening their relevance while making user engagement more real, impulsive, and meaningful. For instance, one of Sony PlayStation's most popular games, "The Last of Us," uses AI to create dramatic performances set in post-apocalyptic America. It gives digital characters the appearance of human intelligence by allowing them to respond to each player through interesting speech and behavior, presented through convincing animation.[26]

This AI reverses the traditional method of developing games. For instance, the producers framed their enemy design around developing adversary characters that seem real enough that players feel bad about killing them. And this type of innovative rethinking can be applied to a range of uses.

AI-driven language processing tools can help people, businesses, and creators to collaborate more effectively and rethink problems of tomorrow.

AI-driven language processing tools can help people, businesses, and creators collaborate more effectively and rethink problems of tomorrow. For example, AI is used by artists and business innovators to make groundbreaking discoveries by studying patterns in the arts and sciences. Meta Foresight demonstrates how emerging AI tools can fire the human imagination and expand access to creative works around the globe.[27]

So for those with the connect and collaborate mindset, customer engagement is much more profound than it is in the previous mindsets, and so requires enterprises to adopt new norms and values, relinquish control of their user communities, and increase the exchange of knowledge with them. This mindset challenges the traditional dominance of in-house research and design (R&D) expertise, instead leading companies to emphasize cocreation processes that result in innovations outside the house, with crowdsourcing fueled by customers' inherent appreciation of being able to help to deliver what they value.[28]

The Empower and Engage Mindset

With its focus on people, planet, and profit, this mindset is often driven by a higher purpose of increasing human and societal involvement.[29] Stakeholders increasingly call for this type of business engagement, so we elaborate upon it in the hope of inspiring leaders to explore the potential of AI in higher pursuits.

When enterprises make an effort to empower their customers and use AI to help them make smarter and better choices, they encourage customer advocacy while bringing new and better opportunities for people and enterprises to redefine social relations and take a proactive approach to pressing societal issues. Starting from a disruptive business-*for*-customer approach, leaders also use AI to solve dilemmas and explore sustainable solutions so they can avoid causing long-term social and environmental harm.

Tesla, for example, has revolutionized the transportation market with its AI-supported, self-driving electric cars, making mobility more efficient and sustainable. The vehicles process a wealth of real-time data from cameras, using computer vision to inform their full autonomy.[30] Tesla uses the PyTorch framework, originally developed by Facebook's AI Research (FAIR)

group, for training and other support tasks such as automated workflow scheduler, model threshold calibration, and passive tests simulation.[31]

AI can also empower customers to engage with sociopolitical issues.[32] The artist Stephanie Dinkins uses AI and media tools to foster conversations about race, gender, aging, and history.[33]

AI can also empower customers to engage with sociopolitical issues.

Enterprises can use augmented reality (AR) and virtual reality (VR) to dramatically increase educational training and entertainment across a variety of instances. Meta's Oculus 2 use spatial AI applications for consumers to explore the hiding place of Anne Frank and her family or to learn how to play the piano or to dance with a robot instructor.[34] Thus, the development of AI, AR, and VR may have broad social, political, and cultural implications for customer engagement using both digital assistants and physical robots.

Lemonade is the first peer-to-peer insurance carrier, and it is scaling up in different areas of insurance, including car, homeowners, pets, and life.[35] This platform enterprise[36] is pairing transferrable AI with cognitive and behavioral psychology[37] to continuously improve customer engagement.[38]

The company's AI is the foundation of its real-time engagement with millennial consumers, but it is Lemonade's overarching empower and engage approach that is reversing the traditional insurance industry model. Meanwhile, its unique combinations of mindset with AI fuel its value propositions and boost its market position. The previous mindsets, by

contrast, focus primarily on tactical uses of AI to optimize logistics, lower costs, offer targeted solutions, and improve customer engagement.

Lemonade's AI-assisted risk assessment minimizes human errors, helping to customize insurance plans, allowing customers to pay only for what they need. The company has also built a transparent fee model based on trust, lower costs, fast claim settlement, and doing social good through peer-to-peer insurance and partial profit sharing.

Lemonade designed its three main bots to be playful, ease customer interactions, and offer real-time services for house renters. Lemonade's *AI Maya* is a conversational virtual assistant that collects information, gives quotes, and handles payments. *CX.AI* answers customers' questions, and *AI Jim* is the claims bot that handles about 30 percent of claims and does much of the work before passing the case on to a human.

Lemonade has also built a forensic graph, which uses AI and behavioral economics to predict, detect, and block fraud. By maintaining a direct relationship with the customer, rather than working through an agent, Lemonade has created a very different business model that boosts its ability to innovate.

Using AI and machine learning[41] to settle claims considerably faster disrupts the often slow and reluctant practices of the traditional insurance industry.[42] And Lemonade's Giveback charity program involves customers directly by allowing them to choose which charities receive each year's unclaimed money.[43] Customers find the company's designed customization services, fast processing, and easy navigation empowering and engaging, as well as entertaining.

As customer engagement preferences become just as important to future innovations as profit and tech, today's leaders must learn to think and act differently.

Developing the Right Strategic Mindset

It is of critical importance that enterprises be able to deploy AI to meet their customers' expectations better than their competitors. The goals of leaders and the identities and needs of their customers determine which AI rationale is right.

However, many enterprises have large portfolios of products and services, so they serve different customer segments with a variety of preferences. In such cases, managers must match their strategic mindset to the customer segments they're targeting, knowing that customer needs will probably change over time and could even conflict in some situations. When customers consider price and environmental impact to be equally important, for example, managers need to continually adjust their strategies and mindsets to succeed in fast-changing and competitive markets.

Table 5 can guide leaders in taking their first steps by showing various ways to apply AI to create relevant, real-time customer engagement. The four business/customer categories show the different outcomes of AI engagement for businesses (e.g., real-time operations) and for customers (e.g., customer convenience) and how to achieve these results with AI.

The future benefits of AI depend more on the underlying mindset than on AI technologies per se. Two of the authors have analyzed

Table 3: Business-*with*-Customer AI Engagement

REAL-TIME CONNECTIONS	CUSTOMER CREATIVITY
· **Stimulate global real-time collaboration and cocreation:** Create dynamic learning and interactive AI platforms. · **Test the functions and add to product and service innovation:** Prototype testing of entertainment platforms.	· **Participate in cocreation processes and application testing:** Return value to customers in terms of improved services and products. · **Creative partnerships to strengthen business-customer relations:** Fast collection and processing of open-source, interactive, user-controlled, online experience and knowledge.

Table 4: Business-*for*-Customer AI Engagement

REAL-TIME ENGAGEMENT	CUSTOMER EMPOWERMENT
· **Solve dilemmas and explore sustainable solutions:** Use AI to avoid long-term negative social and environmental impact. · **Show immediate action in ethical and political issues:** AI can help identify and integrate the interests of stakeholders and higher purposes for the individual and community well-being. · **Serve engaging political causes:** Use AI to create a sustainable business model that satisfies community needs and supports responsible market development.[44]	· **Enable customers to make smarter and better decisions faster:** Use AI to improve the conditions for customers as well as their networks for impactful responsibility-taking. · **Educate customers to critically reflect and to increase self-dependency:** Use AI to help people manage local and global issues on their own. · **Ease customers emotionally and practically:** Increase their self-dependence and help them make a positive difference to themselves and others.

the mindsets of more than 1,000 leaders in many companies and industries, using the information to facilitate the digital transformation processes. The following guidelines are rooted in what they learned, and will help leaders to define their own reasons, their why, for using AI in customer engagement.

Identify the mindset for effective customer engagement with AI that best fits you – and be true to your why. AI offers opportunities across a broad range of company processes as well as market interactions. But to take full advantage of AI, it is essential to ascertain the customer's preferred type of interaction, be true to it, and align your mindset. Customers come with different needs, so tracking customer sentiment, preferences of engagement, and experiences in real time will help enterprises to manage potential frictions that would put customers off.

Meaningful engagement depends on context and time so customers' expectations should be measured,

Employees and leaders who share a mindset will find clearer, more concise ways to use AI, ensuring that their investments succeed.

communicated, and met. However, replacing human services with AI does not necessarily improve customer value or business revenue; it could instead result in bad service experiences. Not all consumers want AI engagement, so it is essential to find a balance between mindset, industry, and customer expectations. Be true to your why and don't call customers kings while treating them as profit-generating servants.

Marketing guidelines in general recommend the connect and collaborate mindset and the empower and engage mindset, but a significant number of businesses still operate under either promote and sell or listen and learn. The classical marketing

orientations are persistent, but many small and medium sized enterprises (SMEs) cannot incorporate customer relationship management (CRM) systems or the AI algorithms that would follow up intelligently. Likewise, many categories of products, such as fast-moving consumer goods, do not require extensive consumer engagement.

Develop real-time capabilities within the enterprise that bring value to the customers. In optimizing their dynamic AI, managers can cultivate a sensitivity to how each mindset functions in real time, adding value for their companies by matching a given mindset with their consumers' preferences. To fully integrate their AI processes, enterprises need to bridge silos, democratize access to data, retrain employees, and identify customer needs.

Increasingly, customers demand that enterprises offer more personal and targeted engagement as well as corporate citizenship.[45] As AI applications push enterprises into

Table 5: AI-Mindset Combinations for Real-time Customer Engagement

Business-*to*-Customer AI Engagement		Business-*with*-Customer AI Engagement	
REAL-TIME OPERATIONS	**CUSTOMER CONVENIENCE**	**REAL-TIME CONNECTIONS**	**CUSTOMER CREATIVITY**
• Optimize production, distribution and sales. • Reach markets rapidly. • Increase product and service reliability.	• Faster and cheaper products and services. • Available offers and communication. • Less friction frees personal resources.	• Stimulate global real-time collaboration and cocreation. • Test the functions and add to product and service innovation.	• Participate in cocreation processes or application testing. • Creative partnerships to strengthen business-customer relations.
Business-*from*-Customer AI Engagement		Business-*for*-Customer AI Engagement	
REAL-TIME RESPONSES	**CUSTOMER RELEVANCE**	**REAL-TIME ACTIONS**	**Customer EMPOWERMENT**
• Instant personal customer data insights. • Optimize and segment processes to fit customer preferences. • Integrate immediate customer ratings.	• Customize digital products, services, and recommendations. • Individualize engagement to strengthen customer relationship. • Fast feedback from customer input increases relevance.	• Solve dilemmas and explore sustainable solutions. • Show immediate action in ethical and political issues such as peace, health, welfare, climate. • Serve engaging political causes.	• Enable customers to make smarter and better decisions faster. • Educate customers to critically reflect and to increase self-dependency. • Ease customers emotionally and practically.

faster marketing cycles and shorter reaction times, both corporations and individuals are under immense pressure, making them increasingly difficult to manage.

The impact of AI on customer engagement may also depend on how well the AI tools, like ChatGPT, work. As it becomes more prevalent, AI, like other enhancement technologies, customers may soon come to take it for granted, making it an essential, rather than a competitive edge.

Leaders should therefore reflect critically on which mindset is most opportune and relevant for effective real-time engagement with their customers, taking into account what it would take to foster that mindset at all points of customer interaction with the enterprise. **Keep the right strategic foundation for AI.** While AI provides new opportunities to optimize customer engagement and value creation, it can easily become a nuisance if it oversteps moral and ethical boundaries. IT investments are often wasted, in part because managers do not identify and share the optimal mindset for customer engagement.[46] It is therefore vital that enterprises explore how AI can best support the chosen mindset before they invest. Employees and leaders who share a mindset will find clearer and more concise ways to use AI, ensuring that the investments pay off. When enterprises rely on leaders' existing mindsets to apply AI, they risk failing to develop the AI solutions preferred by their customers.[47]

Again, the strategic mindsets presented here are not necessarily mutually exclusive. Managers can combine them to build successful AI-enabled customer engagement. But remember that leadership teams with multiple strategic mindsets can create internal inefficiencies and may be seen as opportunistic by customers. It was this kind of inconsistency that gave rise to the concept of greenwashing.

While enterprises might wish to embrace multiple mindsets during a transitional phase, it would be an inefficient strategy for the long haul. Once leaders become aware of their latent mindset, they can effectively examine their own logic and reasoning which in turn will allow them to evaluate how they can best use AI to optimize customer engagement. This analysis will then allow them to assess the relevance and value of AI to their enterprise's particular situation.[48] In short, mindset awareness fuels better leadership for companies hoping to use AI to gain a competitive advantage through effective customer engagement. ■

Author Bios

Pernille Rydén is Dean of Education at the IT University of Copenhagen, Denmark. She earned her PhD in strategic cognition of new technologies from Copenhagen Business School. Her research investigates the influence of strategic sensemaking and decision-making involving business-consumer engagement using digital technologies such as social media, big data, and AI as well as real-time management of businesses and brands in digital business platform ecosystems. pryd@itu.dk

Torsten Ringberg is a professor of marketing at Copenhagen Business School, Denmark. He earned a PhD from The Pennsylvania State University. His interests include sociocognitive theories related to how mental and cultural models and embodied cognition influence perception and understanding, and how consumers and managers create identity. His empirical approaches range from deep-qualitative (ZMET) to experimental methods. He investigates topics such as service recovery, B2B and B2C decision making, and the (ab)use of digital technologies. tri.marktg@cbs.dk

Omar A. El Sawy is Kenneth King Stonier Chair in Business Administration and a professor of information systems in the Data Sciences and Operations Department at the Marshall School of Business, University of Southern California. He earned a PhD from Stanford Business School. He is a fellow of the Association of Information Systems. He is interested in digital business strategy in messy environments, real-time management, the management of AI, and digital platform business models. elsawy@marshall.usc.edu

Endnotes

1. Davenport, T.H. 2020. Enterprise Adoption and Management of Artificial Intelligence, *Management and Business Review.*
2. https://dataprot.net/statistics/ai-statistics/
3. It elaborates on the findings by Rydén, P., Ringberg, T. and Wilke, R. 2015. "How Managers' shared mental models of business–customer interactions create different Sensemaking of social media," *Journal of Interactive Marketing*, 31: 1-16.
4. March J.G. and Simon, H.A. 1958. *Organizations*, NY: Wiley.
5. Gosling, J. and Mintzberg, H. 2003. "The five minds of a manager," *Harvard Business Review, Nov.* 81/11: 54-63.
6. Weick, K. 1995. *Sensemaking in Organizations.* Thousand Oaks, CA: Sage Publications.
7. Stigliani, I. and Ravasi D. 2012. "Organizing Thoughts and Connecting Brains: Material Practices and the Transition. From Individual to Group-Level Prospective Sensemaking," *Academy of Management Journal*, 55(5):1232-1259.
8. Illustration from Rydén, P. Ringberg, T. and Jacobsen, P.Ø. 2017. *Disrupt Your Mindset to Transform Your Business with Big* Data: *A Guide to Strategic Thinking*, Copenhagen, Efficiens.
9. Harris, G.E. 2007. "Sidney Levy: Challenging the Philosophical Assumptions of Marketing," *Journal of Macromarketing* 27: 7-14. and http://www.mckinsey.com/business-functions/marketing-and-sales/our-insights/looking-beyond-technology-to-drive-sales-operations.
10. Fowler, G. 2019. How AI Can Enable Innovation in Sales. https://www.forbes.com/sites/forbesbusinessdevelopmentcouncil/2019/07/09/how-ai-can-enable-innovation-in-sales/; Artfully Intelligent: How AI Helps Customers Shop Smart During Cyber Week: https://corporate.zalando.com/en/newsroom/en/news-stories/artfully-intelligent-how-ai-helps-customers-shop-smart-during-cyber-week
11. https://www.nytimes.com/2022/07/26/business/alphabet-google-earnings.html
12. Pigni, F., Piccoli, G. and Watson, R. 2016. "Digital Data Streams: Creating Value from the Real-Time Flow of Big Data". *California Management Review* 58(3): 5-25.
13. Loftís, L. 2017. "Maximizing Moments of Truth: Creating Meaningful Real-Time Customer Interactions". SAS Best Practices, SAS Institute.
14. cf. https://www.fool.com/investing/2022/03/30/automation-reducing-sales-agent-jobs-lemonade/
15. https://emerj.com/ai-sector-overviews/artificial-intelligence-at-amazon/
16. https://corporate.walmart.com/newsroom/2021/06/24/headline-how-walmart-is-using-a-i-to-make-smarter-substitutions-in-online-grocery-orders
17. https://www.mycustomer.com/experience/engagement/five-examples-of-companies-using-ai-to-improve-customer-experience; https://www.businessinsider.com/ap-macys-tests-artificial-intelligence-tool-to-improve-service-2016-7?r=US&IR=T
18. Prevalent in high-turbulence tech-driven industries. Eisenmann T.R. and Bower, J.L. 2000. "The entrepreneurial M-form: Strategic integration in global media firms," *Organization Science*, 11/3: 348-355.
19. Loftís, L. 2017. Maximizing Moments of Truth: Creating Meaningful Real-Time Customer Interactions. SAS Best Practices, SAS Institute.
20. Syam, N. and Sharma, A. 2018. "Waiting for a sales renaissance in the fourth industrial revolution: Machine learning and artificial intelligence in sales research and practice," *Industrial Marketing Management,* 69(2): 135-146.
21. Rydén, P. and El Sawy, O. A. 2019. "Real-Time Management in the Digital Economy," in T. K. Das (ed.) *Time Issues in Strategy and Organization,* Information Age Publishing, pp. 59-92.
22. Zaki, M, McColl-Kennedy, J. R, and Neely A. 2021. Using AI to Track How Customers Feel - In Real Time, May 04. https://hbr.org/2021/05/using-ai-to-track-how-customers-feel-in-real-time.
23. https://www.globenewswire.com/fr/news-release/2021/09/14/2296756/0/en/Kore-ai-Introduces-BankAssist-To-Empower-Customers-With-Extraordinary-Self-Service-Experiences.html
24. Edelman D.C. and Abraham, M. 2022. Customer Experience in the Age of AI. *Harvard Business Review* (March–April 2022). https://hbr.org/2022/03/customer-experience-in-the-age-of-ai
25. https://www.pointillist.com/customer-journey-management-software/
26. Botta, M. 2015. Infected AI in The Last of Us, *Game AI Pro* (2), Chapter 33. https://www.gamedeveloper.com/design/endure-and-survive-the-ai-of-the-last-of-us
27. How AI Will Elevate the Future of Creativity. https://www.facebook.com/business/news/insights/how-ai-will-elevate-the-future-of-creativity
28. Zahra, S.A. and Nambisan, S. 2012. "Entrepreneurship and strategic thinking in business ecosystems," *Business Horizons,* 55(3): 219-229.
29. Like Kramer and Porter's Creating Shared Value (CSV) in Porter, M. and Kramer, M. 1999. "Philanthropy's New Agenda: Creating Value," *Harvard Business Review* 77(6): 121-130.
30. Light Detection and Ranging (LIDAR) is a remote sensing method that uses light in the form of a pulsed laser to measure ranges (variable distances) to the Earth. These light pulses, combined with other data recorded by the airborne system, generate precise, three-dimensional information about the shape of the Earth and its surface characteristics.
31. https://www.fireblazeaischool.in/blogs/how-tesla-uses-ai-and-cv/
32. du Boulay, B., Poulovassilis, A., Holmes, W. and Mavrikis, M. 2018. What does the research say about how artificial intelligence and big data can close the achievement gap? R. Luckin (ed.), *Enhancing Learning and Teaching with Technology*. London, Institute of Education Press, pp. 316–327.
33. Dinkins, S. 2020. When Words Fail: https://www.onassis.org/whats-on/festival-you-and-ai-through-the-algorithmic-

lens/artworks/whenwordsfail?fbclid=IwA R3wFJR2zEdyNOMyDmU5qWQ37RETel8JCc OI9D7mNl2sr2Ro9ln00tvLcigwhenwordsf ail?fbclid=IwAR3wFJR2zEdyNOMyDmU5q WQ37RETel8JCcOI9D7mNl2sr2Ro9ln00tv-Lcigwhenwordsfail?fbclid=IwAR3wFJR2z EdyNOMyDmU5qWQ37RETel8JCcOI9D7m Nl2sr2Ro9ln00tvLcig

34. https://ai.facebook.com/blog/powered-by-ai-oculus-insight/

35. Established in 2016. Alois, J. D. 2017. *Lemonade Updates on Platform Growth. Adds 14,300 Customers in 8 Months,* Crowdfund Insider.com, June 1, 2017, www. crowdfundinsider.com/2017/06/101283-lemonade-updates-platform-growth-adds-14300-customers-8-months/.

36. Insurtech refers to technological innovations that improve the efficiency of the insurance industry cf. https://www.pwc.com/gx/en/industries/financial-services/fintech-survey/insurtech.html

37. https://www.businessinsider.com/dan-ariely-joins-lemonade-insurance-2016-2?r=US&IR=

38. Insurance video: https://www.youtube.com/watch?v=flSLI2JmWVE.

39. We do not use this case example to endorse Lemonade nor belittle its achievements. That its stock has been slumping since March 2021 and has reported net loss does not detract from our purpose of illustration, cf. https://finance.yahoo.com/news/lemonade-down-75-disruptive-growth-172008105.html

40. Lemonade employs around 500 people for tasks that cannot be automated.

https://rickhuckstep.com/lemonade-5yr-anniversary/

41. https://digital.hbs.edu/platform-rctom/submission/lemonade-reinvents-the-insurance-industry-with-machine-learning/

42. https://emerj.com/ai-sector-overviews/artificial-intelligence-in-insurance-trends/

43. Data scientists have modeled an optimal mix of internal and external 'reinsurance'. The amount available for Giveback will average less than 40% due to claims. https://www.mymoneyblog.com/lemonade-vs-mutual-insurance.html

44. Boothroyd 1991 (p. 130) calls this "the creation of a new community in 'non-gemeinschaft' form that aims to fulfill basic human needs," Boothroyd, P. "Community development: The missing link in welfare policy, " In Kirwin B. (Ed.), *Ideology, development and social welfare: Canadian perspectives*, Toronto: Canadian Scholars Press.

45. Feraj, S., Pachidi, S., and Sayeg, K. 2018. "Working and organizing in the age of the learning algorithm," *Information and Organization* 28(1): 62-70.

46. Rydén, P. Ringberg, T. and Jacobsen, P.Ø. 2017. *Disrupt Your Mindset to Transform Your Business with Big Data: A Guide to Strategic Thinking,* Copenhagen, Efficiens

47. Tenner, E. 2018. *The Efficiency Paradox, What Big Data Can't Do*, New York, Alfred A Knopf.

48. Kaplan, S. 2011. "Research in cognition and strategy: Reflections on two decades of progress and look to the future," *Journal of Management Studies,* 48(3): 665-695.

Book Review: Balancing Human and AI Control to Achieve Meaningful Customer Engagement

Human-Centered AI
by Ben Shneiderman

Avi Parush

Israel Institute of Technology

The introduction of smart systems, automation, and autonomy is strong evidence of the proliferation of artificial intelligence (AI) technologies. In every facet of our lives, AI is increasingly present through various service systems including targeted, personalized, and anticipatory marketing techniques rooted in big data analytics, chatbots that provide automated service, sometimes at critical touchpoints with customers, and more.

We assume that AI technologies within service systems will increase the customer's engagement with the service itself and the company providing that service. Yet, paradoxically, one of the biggest concerns about the proliferation of AI is that people will gradually lose control, become less involved, and disengaged. Is AI's use for human engagement a double-edged sword?

It is in scrutinizing this question that I value the timely significance of the book, *Human-Centered AI*, by Professor Ben Shneiderman from the University of Maryland.[1] It voices a critical call for AI to be human-centered and provides a constructive and practical guide on how to go about it. The book is organized logically: starting with AI's fundamental concepts and its philosophical and practical underpinnings, introducing Shneiderman's human-centered AI (HCAI) framework, and moving on to the implications of AI model design, a discussion of AI governance structures, and a future agenda.

The book offers insightful lessons and practical takeaways, too many for me to cover in this short review. I would like instead to focus on what I consider to be the key aspects and most important ideas: the HCAI framework and its broader implications, especially in the

context of AI for customer engagement. I offer an interpretative commentary, rather than a traditional review, which revolves around several facets of human engagement with AI as a way to view Human-Centered AI.

HCD, HCI, and HCAI

When I first heard about human-centered AI, I had a feeling of déjà vu. I was reminded of the birth and evolution of human-centered design (HCD) in human-computer interaction (HCI). Shneiderman has been one of the torchbearers for incorporating the human in the HCI equation. Is HCAI then a natural evolution of HCD, or are we in need of a paradigm shift? How disruptive will AI be when it comes to tackling the challenges of taking a human-centered approach to the design and development of any technology that affects people?

Shneiderman's HCAI framework and its broader implications suggest that the transition to human-centeredness in AI requires a more fundamental change than the movement towards HCD.

Balanced engagement

Quite early in the book, the author mentions 'human control' in HCI. I was immediately reminded of a series of public debates on user interface design between professors Shneiderman and Pattie Maes about a quarter century ago. Shneiderman has been a long-time proponent of direct manipulation, or giving people control and predictability in user interfaces, and thus more human control in HCI, while Maes was in favor of giving software agents a larger role in the interaction – that is, more autonomy to act on the user's behalf given the ever-increasing proliferation of choices.[2]

The HCAI framework reframes our thinking about human control vs. computer automation. One of the key points Shneiderman makes

Figure 1: Shneiderman's Two-dimensional HCAI Framework, (Reproduced from the book with permission.)

Figure 1: Designers need to prevent the failures from excessive automation and excessive human control (gray areas).
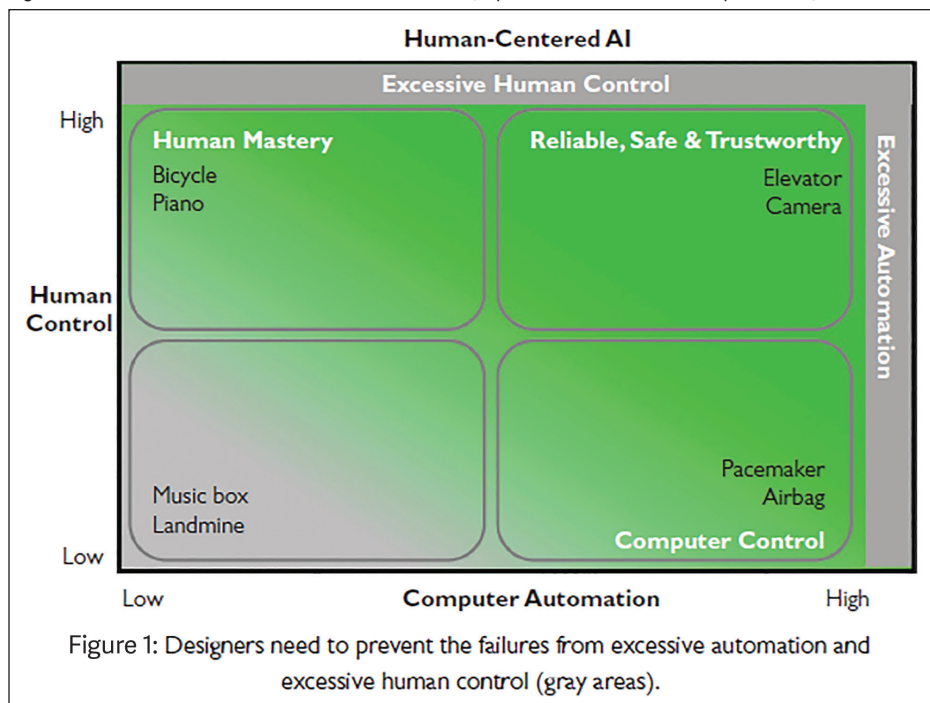
in this book is that more control to one does not necessarily mean less control to the other; this is not a zero-sum game. It is also not the unidimensional scale of automation and autonomy that commonly appears in contemporary thinking about the evolution of AI and intelligent systems, which holds that there are certain tasks only humans are good at and others that computers can do better. The lines have blurred. Tasks that typically fall in the realm of humans, like medical diagnoses and driving a car, are increasingly being automated.

In a way, the book is a call to shift away from the good old MABA-MABA (men-are-better-at vs. machines-are-better-at) way of thinking and to reframe the allocation of functions in human-AI systems.

The HCAI framework suggests a more balanced, win-win view of the degree of control and engagement by humans and AI software in the design, implementation, and use of intelligent systems. The book depicts the problem of human and

software automation and autonomy as two-dimensional: vertical for low to high levels of human control, and horizontal for low to high levels of computer automation.

The HCAI framework calls for an appropriate balance between human control and computer automation. The idea is to design the system in such a way that it strikes a balance between human capabilities and software automation and autonomy.

The framework also suggests that the balance is not static and prescriptive, but rather dynamic and adaptive. It depends on a variety of factors and circumstances, such as human and technology capabilities and the need for more – or less – active human involvement and engagement. The implication of such a balance for AI-based customer engagement is that it should be bidirectional.

It is not only the AI that drives the engagement; it is also the extent of human involvement and control, and specifically meaningful control. The concept of meaningful

Figure 2: A Proposal Adapted from Shneiderman's HCAI Framework
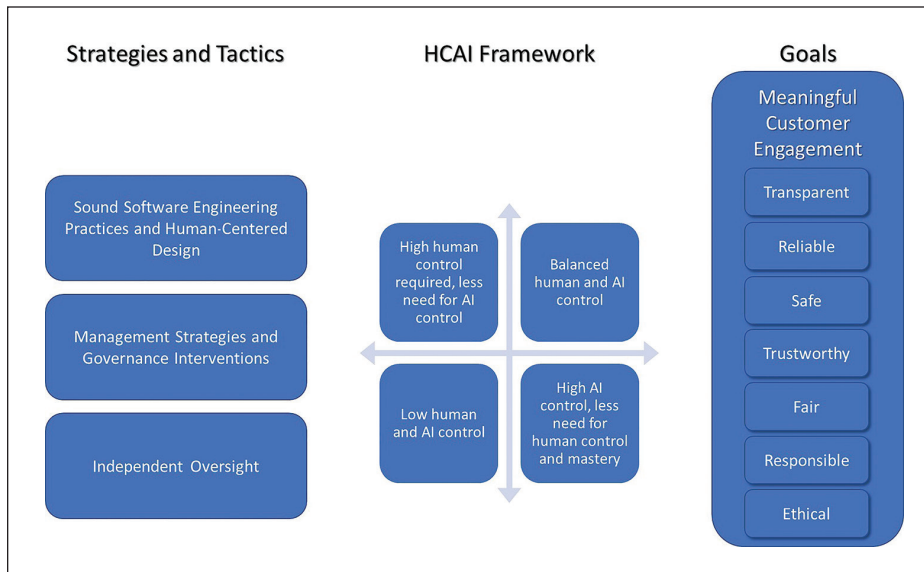
human control is mentioned several times in the book and often is used in the context of lethal autonomous robots. Through it, we can talk about meaningful customer engagement, which means shifting some of the control to the customer through greater customer awareness, understanding, and influence over an AI-driven engagement.

Multilevel and continuous engagement

Human-Centered AI is a call to go beyond HCD approaches to design and testing, to go beyond considering only the end user, and beyond the specific context of the usage or the presence of AI. It urges us to go beyond user experience and usability, and even beyond designing explainable AI (XAI), to make it more human-centered.

Indeed the book suggests that the implications of this question are far wider and potentially more critical than those of our struggle for human-centeredness in HCI. The author urges us to consider reliability, safety, trustworthiness, fairness, and ethics, as well as organizational and governance issues. He calls for multilevel and

continuous human involvement and engagement in the entire lifecycle of the AI system.

This involvement starts with individuals and extends to organizations, regulatory bodies, agencies, nations, and societies. Involvement and engagement are not only incorporated in the design, development, and deployment of the AI, but also run throughout the lifecycle of the system, embodied in governance structures that emphasize organizational and business aspects, regulation, and oversight. Human-centeredness in AI is an ephemeral constant: It should always be there, yet it is ever changing.

> There is a lot of interest in 'responsible AI,' which is good news. The disconcerting news is that there is far less interest in 'safe AI.'

Multi-faceted engagement

The human-centeredness of AI, according to Shneiderman, is also about the reliability, trustworthiness, fair-

ness, responsibility, ethics, and safety of AI systems. All of these are tightly linked. Yet, a quick look at online search trends at the time of writing this review reveals this: There is a lot of interest in 'responsible AI,' which is good news. The disconcerting news is that there is far less interest in 'safe AI.'

I would like to pause here and dwell for a moment on this neglected safety question. Safe AI is becoming a critical challenge. Mica Endsley, the former U.S. Air Force chief scientist, has said that "autonomy systems that drive vehicles (whether fully autonomous or not) should have to pass driving tests just like people do."[3]

Shneiderman's emphasis on safety and, importantly, safety within organizations, their business strategies, and their culture, is a significant and timely component of the shift needed in our thinking about humans and AI. In AI-driven customer engagement, such as customer targeting and personalization, it is critical to mitigate errors, failures, and accidents.

Implementing human-centered AI can ensure a relevant transparency through which customers can recognize and understand system errors and failures, like irrelevant targeting or a privacy breach.

HCAI future research agenda

In both Part 3 (Design Metaphors) and Part 5 (Where do We Go from Here?), Shneiderman proposes scientific and practical ways to use the key messages of the book as we move forward. One important way forward is to focus on the challenge of metrics. What are the metrics for better AI systems? Shneiderman addresses such challenges particularly with respect to assessing trustworthiness of AI systems (chapter 25 in Part 5).

His message is that any future scientific research and applied innovation will have to use valid and

reliable metrics to measure and assess the extent to which the AI system is human-centered. This in turn will allow us to implement adequate governance structures, ensuring that AI systems are reliable, trustworthy, ethical, safe, and human-centered.

> In AI-driven customer interactions, the degree of meaningful customer engagement is a potentially important metric.

In AI-driven customer interactions, the degree of meaningful customer engagement is a potentially important metric. In the spirit of human-centered AI and balanced engagement, this metric could assess the impact of engagement in which some control is shifted to the customer and then balanced by some AI control, all while ensuring that the organization meets its strategic goals.

Closing and practical takeaway

I started this commentary with the following conundrum: AI can make people gradually lose control, becoming less involved and less engaged. Is AI, then, a double-edged sword with regard to human engagement?

Shneiderman's HCAI proposal implies that the use of AI in customer engagement should adhere to an appropriate balance between human and automated control so as to drive meaningful customer engagement. I propose adjusting the HCAI approach, in keeping with the schematic diagram above.

The ultimate goal for HCAI in building meaningful customer engagement is that it be transparent, reliable, safe, trustworthy, fair, responsible, and ethical. To achieve these goals, consider my adaptation of the HCAI framework for a balanced engagement.

> The ultimate goal for HCAI in building meaningful customer engagement is that it be transparent, reliable, safe, trustworthy, fair, responsible, and ethical.

Creating human engagement with systems in general, and specifically with service systems, is mostly about drawing people in and keeping them loyal and engaged.

In *Human-Centered AI,* Shneiderman suggests strategies and tactics for implementing the HCAI framework. I find several of these to be particularly relevant to AI-driven customer engagement. These include sound software engineering and human-centered design methods; appropriate governance and management, including timely and proportionate interventions, and independent oversight throughout the system's lifecycle. These measures help managers to enact the HCAI framework and achieve their goals.

Shneiderman's HCAI argues that human engagement with systems should not be just a result of employing AI technologies, but an inherent part of the system. Human engagement should be built into AI's conception, design, development, and deployment, ensuring that humans play a meaningful and active role in AI's entire lifecycle. ■

Author Bio

Avi Parush earned his PhD in Experimental Psychology from McGill University, Montreal in 1984. He is an associate professor in the faculty of data and decision sciences at The Technion, Israel as well as an emeritus professor at Carleton University, Ottowa and an adjunct professor at Queen's University, Kingston. His current research focuses on teamwork in complex and critical situations, human factors in healthcare, human-robot interaction, human behavior relative to autonomous vehicles, and conventional and simulation-based training.

Endnotes

1. Shneiderman, B. *Human-Centered AI.* Oxford University Press, 2022.
2. "Direct Manipulation vs. Interface Agents." *Interactions,* pp. 42-61, November & December issue, 1997.
3. Mica Endsley, May 15 2022, LinkedIn

Book Review:

Artificial Intelligence for Sustainable Value Creation, Margherita Pagani and Renaud Champion, eds.

Artificial Intelligence for Sustainable Value Creation is an edited volume whose authors offer a detailed and insightful exploration of both the possibilities and the challenges of widespread use of artificial intelligence (AI). It analyzes the effects of AI on business and society with an eye to assessing what we know about managing other information systems, strategy, and marketing, reexamining this knowledge in the context of AI. The book's contributors explore how human-centric AI systems create value for organizations, discussing three main categories: ethical value, societal value, and business value.

Charles Hofacker
Florida State University

In the first chapter of the edited volume, *Artificial Intelligence for Sustainable Value Creation*, Margherita Pagani and Renaud Champion present the concept of human-centric AI, by which they mean AI systems that perceive, reason, make decisions, and act according to human values. These values include accountability and avoiding bias, in this case algorithmic bias.

The implications for customer engagement are clear: AI must be designed to avoid inadvertent unethical acts. The concept reminds me of the famous 'Paperclip AI' thought experiment: imagine a world in which a powerful AI is trained only to maximize paperclip production, without any other considerations. After a few years the entire surface of the planet would be miles deep in paper clips. So much for human existence!

We have already seen, and so do not need to imagine, what happens when algorithms devoid of human ethics are tuned to optimize consumer engagement. Such algorithms have succeeded in producing massive engagement, but they have also harmed societal welfare and fomented polarization and alienation.[1]
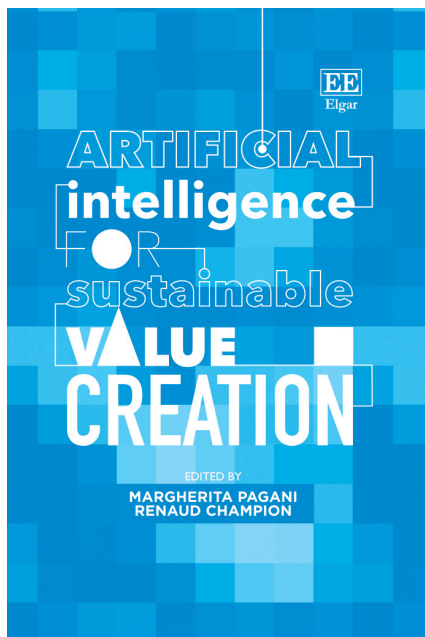
Making AI human-centric, as these authors insist we must, would make a promising start for any AI designed to foster consumer engagement. If you are a science fiction fan, you may by now have been reminded of Asimov's First Law of Robotics, "A robot may not injure a human being or, through inaction, allow a human being to come to harm." Similarly, an AI trained for consumer engagement should follow the dictum "first, do no harm."

In Chapter 2, Yihyun Lim takes a similar Asimovian tack, contemplating a values-driven design for AI. Lim notes that human values need to be deeply embedded in design, which is not so easy, because these values must be carefully elicited from stakeholders. What values are precious to those who act and engage on behalf of the brand or with respect to it?

One implication here is obvious: few consumers hope or intend to become outraged and polarized in service to some brand's pursuit of profit.[2] Instead, Lim suggests that an AI designed to increase engagement or other organizational goals should do so with an eye to such values as protection, acceptance, assistance, and acknowledgement.

Platforms vis-à-vis other business models

Chapter 3 addresses more ordinary business concerns as Omar El Sawy, Milan Miric and Margherita Pagani point out that AI generally exists in the context of some platform business model. This context is especially important because customer engagement tends to take place on a platform. Just as the rich are not like you and me, platforms are clearly not like other business models.

The authors go on to enumerate the ways in which they are different including network effects, in which more users equal more value; a community focus; and the need for governance mechanisms. They also point out that, in some cases such as network effects that accelerate data collection, engagement platforms and AI actually complement each other.

What remains less clear, however, is how factors like the need for governance relate to the AI imperative of maximizing efficiency, utility, and predictability. El Sawy et al. propose configurational theories as a way of understanding the dynamics, but it is unclear how these might enlarge our understanding of how AI and platform governance could coexist, unless it is at the expense of one over the other.

And if governance is important, which it surely is, should we insist that AI design adhere to the principles of common governance? Put another way, if I am going to have my attention engaged by an AI more clever than I, should I not have some say in this process?

In Chapter 5, Christine Balagué touches upon many of the ethical issues noted above, and adds a concern about the opacity of AI algorithms, their black-boxiness, as it were. It would be interesting to know whether opacity is part of what makes clickbait so engaging. Put this another way, if an AI suggests that I read a bit of content *x* and then carefully describes in detail how it predicted that I would be drawn to *x*, would the explanation make *x* less attractive to me?

Balagué also painstakingly documents the problem of ethnic and racial discrimination. Lacking moral judgment, an AI trained merely to boost engagement would be happy to push hateful material for those who might be drawn to it. This, too, is a platform governance issue. Neither I nor the authors are prepared to offer any easy solutions but conceptualizing it as a governance issue is a good start.[4]

In conclusion, several of the book's contributors suggest that there are important factors that companies should consider as companies go merrily about training their AI models to foster customer engagement. Unfortunately, the consequences of ignoring these factors have already been realized on Twitter and its less prominent brethren.

More generally, we might benefit from thinking of engagement as one of a set of behaviors that generate customer value. In this light, the perils of AI running amok in its efforts to generate engagement are a special case in our broader efforts to optimize the lifetime value of customers.[3]

Artificial Intelligence for Sustainable Value Creation sheds a bright light on how human-centric artificial intelligence could create sustainable value for customers and society. It also suggests practical ways for us to develop ethical customer engagement.

Any way you look at it, it is ever more apparent that the fabled Silicon Valley admonishment to move fast and break things has driven AI to do just that with respect to consumer engagement. This book offers some hope that academics, industry, and enlightened regulators might take the implications seriously and begin to mitigate the breakage all around us. ∎

Author Bio

Charles Hofacker is the Persis E. Rockwood Professor of Marketing at Florida State University's College of Business. His research investigates the intersection of marketing and information technology and has appeared in the *Journal of Marketing Research, Journal of the Academy of Marketing Science, Psychometrika, Management Science, Journal of Advertising Research* and more. Hofacker is the moderator of ELMAR, an electronic newsletter and community platform for academic marketing with more than 8,000 subscribers.

Endnotes

1. Berman, Ron and Zsolt Katona (2020), "Curation Algorithms and Filter Bubbles in Social Networks," Marketing Science, 39 (2), 296-316.
2. Hollebeek, Linda D., David E. Sprott, Valdimar Sigurdsson, and Moira K. Clark (2022), "Social Influence and Stakeholder Engagement Behavior Conformity, Compliance, and Reactance," Psychology & Marketing, 39 (1), 90-100.
3. Libai, Barak, Yakov Bart, Sonja Gensler, Charles F. Hofacker, Andreas Kaplan, Kim Kötterheinrich, and Eike Benjamin Kroll (2020), "Brave New World? On Ai and the Management of Customer Relationships," Journal of Interactive Marketing, 51, 44-56.
4. Zeng, Helen Shuxuan, Brett Danaher, and Michael D. Smith (2022), "Internet Governance through Site Shutdowns: The Impact of Shutting Down Two Major Commercial Sex Advertising Sites," Management Science, 68 (11), 8234-8248.

©shutterstock.com/thinkhubstudio

White Glove Service: AI in Wealth Management Raises Client Engagement at Scale

Brian Lincoln, IBM Consulting

Robert Grant, IBM Consulting

Suresh Iyengar, IBM Consulting

Brian Lincoln, Robert Grant, and Suresh Iyengar explain why firms seeking to win in the growing wealth management market must apply AI for intelligent document processing and how doing so will drive clients' engagement, meeting their experience expectations and scaling operations for efficiency.

The rush of new investors to the markets over the past two years and the growing pool of affluent individuals seeking advice have created a significant strain on wealth management operations.[1] Clients want firms to know them and expect the same digital convenience they experience in other parts of their lives, including instant access to information and insights about their portfolios, digital execution of all transactions from their mobile phones, and immediate or rapid feedback, both digitally and in person.

These expectations challenge wealth managers to meet an ever higher standard of engagement, because even in today's largely digital world, there is no escaping the need for physical documentation and paper. Managing that paper is often frustrating for clients and can also constitute an undue burden on the firm's back and front offices. Worse still, repeatedly keying information into multiple systems generates errors.

But now, leaders have started to apply artificial intelligence (AI) and machine learning (ML) to ease this bottleneck by dramatically improving intelligent document processing (IDP) capabilities to allow for automatic or straight through processing (STP) of client requests, all of which were previously handled manually.

IDP uses natural language processing (NLP), optical character recognition (OCR), and deep learning to read and understand documents such as tax forms, extracting relevant information automatically and at scale. IDP is critical to transforming the wealth management process, building client engagement, and driving revenue growth.

Transforming the Client Experience

Not only can IDP reduce document management costs, it can dramatically transform the client experience, reducing processing times, limiting or eliminating errors, and giving clients regular updates as their requests are processed. The same type of machine learning models that parse document content can also be used for personalization, introducing new capabilities which build progressively into ever more sophisticated uses, all while engaging clients.

Leaders can also use AI capabilities to upgrade high-end experiences, extending the white glove service to digital interactions (mobile, web) so that clients can upload all types of documents, for example in estate planning, and see them processed instantly. This immediate understanding can engage clients with tailored portfolio news, investment recommendations for self-directed and advisory clients, and less paper-intensive estate planning.

Many wealth management use cases are document-intensive, from bringing a new client on board, to administering trusts, to settling the affairs of a deceased client. Each of these operations requires documents from the client that can be categorized into three types: structured, semi-structured, and unstructured.

> The same type of machine learning models that parse document content can also be used for personalization, introducing new capabilities which build progressively into ever more sophisticated uses, all while engaging clients.

Structured forms have fields in consistent patterns and locations on the page, for example, the firm's form for opening a new account. Semi-structured forms, like utility bills, have similar data, but in fields that could be anywhere on the page. Unstructured forms follow no strict pattern and require context and understanding to extract the data, as is the case with a trust agreement.

To complicate matters, computers, like people, can find correctly interpreting handwriting to be very challenging because of variations in document resolution, ink color, and smudges.

Where to Start

The solution to this uncertainty lies in creating technology that can span a client's omnichannel needs while supporting the myriad uses inherent to large firms. Mature firms tend to use complementary technologies arranged into a pipeline, with a variety of models and technologies that extract and deliver the data as well as drawing insights. The pipeline allows multiple technologies and machine learning approaches to cooperate to individually tailor operations and clients' experiences.

Using a modular pipeline approach to IDP allows the system to manage and process documents branch stores, mobile apps, and websites, all while adhering to WORM (write once, read many) record storage requirements, which insure against tampering. As the document is ingested, firms can use a variety of models to classify the document. This classification helps tailor the downstream NLP models to ensure that they achieve very high accuracy while extracting different types of content.

Foundation models, which include large language models, such as OpenAI's GPT-3, Google's LaMDA, T5, BERT, Facebook's RoBERTa, and more, have

Table 1: Transformation Imperatives

Current State	Future CX with Full AI Implementation
Digital experience for clients and advisors • Only basic digital tools available • Lack of digital/advisor tool integration for hybrid advice and servicing	• Hybrid advice delivered seamlessly with personalization and collaboration options • Smart advisor platforms powered by virtual assistants and cognitive client insights
Holistic personalized offers and advice • Product-oriented capabilities, with need to re-enter data across tools • Little personalization	• Enable client 360 across enterprise • Personalized offers, advice and servicing, with real-time insights
Reimagined operations • Manual processing, paper intensive, slow • High NIGO (not in good order) rates and compliance burden	• Harmonized, segment-appropriate self-service • Data reuse and content intelligence enable digital servicing with embedded compliance, drastically reducing NIGOs
Business agility • Lengthy dev / test cycles limit release frequency • Segment siloes hinder upgrades from retail to wealth	• API strategy supports integration to 3rd party ecosystem for new products and services • Modernization enables segment mobility

exceptional NLP and generative AI capabilities. They are pre-trained on a massive corpus of data which allows them to be fine-tuned with a comparatively small amount of data, termed few-shot learning (FSL), to understand documents specific to a particular domain or enterprise. This fine-tuning requires access to domain expertise. For example, a trust lawyer can explain legal terms, like beneficiary or executor, which can be expected to appear in trust documents. The foundation models can thus learn to understand the context and relationships of content from complex unstructured documents.

Managers can use the generative capability of these foundation models to summarize complex documents, enable contextual semantic searches, and even create personalized drafts of communications with clients. AI, in these cases, acts as a copilot for the wealth advisor, instantly processing a large variety of documents to offer suggestions and generate per-sonalized communication in natural language as well as completing business operations tasks.

Small Steps, Big Results

But even comparatively simple approaches using structured forms can produce impressive results. When IDP improves accuracy from 75 percent to 95 percent, it may not seem immediately significant. But firms that adopt the technology find that they can manage exceptions and respond to volume spikes more effectively with one tenth the support staff, which ultimately saves them millions of dollars in support costs, assuming more than five million documents processed per year on average.

For example, a US-based wealth management firm which processes more than 10 million documents a year, mainly having to do with account openings and moving money, saw the following results after applying IDP:[2]

• **Reduction in request cycle time:** Reduced data entry time across many services up to 95 percent (range 30 percent to 95 percent)

• **Reduction in management costs:** Ten times lower document operations and management costs (range 20 percent to 90 percent)

• **Increase in capacity:** Handled three-times the volume when business spiked with no increase in operations staff or expense (range double to triple)

• **Improvement in client engagement:** Improvement of Net Promoter Score (NPS) by 10 to 20 percent across multiple programs. Increased positive client feedback, net new assets relative to peers and comparative transaction volumes (range 10 percent to 30 percent depending on scope of program).

Lessons Learned

Many wealth management firms have conducted proofs-of-concept

to use AI for intelligent processing. Their results have been mixed for a variety of reasons. Here are the best practices for a successful outcome:

- ✓ **Use case suitability:** Start with a use case that is achievable and can deliver short-term business value.
- ✓ **Technical fit:** Ensure that the technology is sufficient to deliver the desired outcomes and will support the broader context.

- ✓ **Establish a pipeline:** Tie several capabilities together to solve for multiple use cases, although there is no one-size-fits-all solution.
- ✓ **Strive for STP:** Determine what you need for straight through processing (STP) of requests and use that as your north star.
- ✓ **Establish enterprise scale:** Most lines of business face the same challenges with document ingestion. Centralizing this capability

across the enterprise can pay off significantly and gain pricing leverage with vendors.
- ✓ **Tooling:** Commercially available tools typically produce better results right out of the box, at scale, and with limited tuning.
- ✓ **Assume it matters:** AI produces the service quality necessary to build engagement. At one private bank, credit card service issues were the leading cause of attrition. For clients, IDP can be the difference between delight and disappointment. ■

Author Bios

Brian Lincoln has more than thirty-five years of experience in wealth and asset management and more than forty years of experience in financial services in industry and consulting. He has led business transformations across all segments of wealth management. Recently, he has been supporting the application of analytics and advanced technology from IBM and third parties to improve client engagement, sales, and advisory effectiveness, risk management and efficiency in wealth and asset management. Brian.lincoln@us.ibm.com

Robert Grant has delivered technology and operations solutions for wealth management and banking firms for over a decade. He is an operations process technology expert and has worked with leading financial firms in the US on global solutions. He is an IBM Master Inventor, with more than ninety patents filed in financial services, process re-engineering, and more. He leads IBM's technology department for the wealth management practice, heading a team of data scientists and technology consultants. RHGrant@us.ibm.com

Suresh Iyengar is Senior Partner and Global Leader for Wealth Management, Custody and Securities Services for IBM Consulting. In leadership roles for more than twenty-five years at fintechs (Jemstep, Intellect Design, Reuters, and BANCS) and banks (Citibank, SCB) alike, Suresh radically transformed business models and grew businesses, driving efficiency through disruptive technologies, by re-inventing wealth management and custody. His experience spans leadership, business transformation, P&L, operations, and advisory roles, including growth initiatives and post-merger and acquisition scenarios. Suresh.iyengar@ibm.com

Endnotes

1. Oliver Wyman, Morgan Stanley Wealth and Asset Management Report 2022.
2. Using a year-to-year comparison, with a volume spike in the COVID era. We provide a results range to reflect comparisons with four other leading wealth management firms.

©shutterstock.com/Andrey Suslov

Alan Schulman
UpperRight

Stacey Lynn Schulman
Hi: Human Insight

Resurrecting Jimi Hendrix:
The Power of AI to Expand Consumer Engagement Through Musical Fan Cultures

Since the earliest days of written and recorded music, many of history's greatest composers, musicians, and performing artists have created their own unique sound, a signature style or approach that is instantly recognizable and identified with them. Most of these legendary artists also left behind unfinished works or died far too young, leaving us to wonder what they might have produced had they continued to develop and expand their craft.

In recent years, technologists have tasked artificial intelligence and machine learning (AI/ML) with replicating, reimagining, and expanding upon the work of these historic artists – how might their sound have evolved had they lived on?

AI is already creating content, particularly imagery and text. However, people have only recently begun to explore generative music through AI-generated machine learning. With this in mind, we set out

Alan Schulman and Stacey Lynn Schulman describe how generative AI can make classic, iconic musical styles feel new again, helping people to engage with artists who are gone, but not forgotten. This technology could fuel consumer engagement with brands that invest in musical cultures and branding.

to assess how AI-generated works in the style of iconic musicians would engage consumers. With so many brands trying to establish relevance with youth culture, expanding the musical library of icons among engaged fans could create new opportunities to ignite consumers' passion across multiple life stages.

We explored how listeners react to AI-generated music that endeavors to interpret and advance the signature sound of several world-renowned artists. By analyzing widely available consumer sentiment as well as data drawn from our independent research, we explored the delicate boundaries between art and technology and designed a framework for assessing how AI-generated music can achieve the aesthetic and commercial saliency that consumers expect from human musicians. Our findings will have a broad impact on consumer engagement with brands that invest in associating with musical fan cultures and sonic branding.

AI, creativity, and the human element

Assessing AI for its propensity to engage consumers is complicated in the artistic realm. Art and its value are notoriously debated by human critics, audiences, and creators. In our effort to develop a framework on which to evaluate AI-generated music, we drew upon this fertile ground, unearthing the characteristics that underpin most artistic debates. Whether comparing fine art vs. comic books, rock music vs. pop, or the works of different periods of an artist's development, four characteristics come up repeatedly: authenticity, emotional vibrance, experiential/experimental triggers of creativity, and engagement with the audience. These concepts are woven together such that the mere act of creation and consumption of art are seen as uniquely defining the human condition.

Artificial intelligence quite simply fails the sniff test when applied to reproducing the essence of our humanness.

With respect to creative pursuits, *artificial* intelligence suffers from an unfortunate moniker. AI quite simply fails the sniff test when applied to what is considered the essence of our humanness. In a world that is becoming increasingly automated, creativity is the safehouse humans return to in defense of their value and unique contributions to the universe. Whether you believe creativity is defined by the end-product or by the process of creation itself,[1] many believe that the value to both the creator and the audience is in the representation and evocation of something machines will never be: deeply emotional and even irrational.

Humans are biased evaluators of AI's artistic contributions largely because of the importance we ascribe to emotion in the act of creation. If algorithms do not feel, can machines truly create?

Moreover, will consumers engage with and value AI creations as they have human ones?

Selecting AI-generated music to study

We examined two different and discrete sets of AI-generated music encompassing various renowned composers, musicians, and musical genres. These compositions were part of two distinct projects using different AI technologies. Both projects are publicly available online:

The Lost Tapes of the 27 Club[2] is an AI project developed by Over the Bridge, a Toronto-based organization that set out to raise awareness about mental health issues in the music community, a group that has struggled with mental health at a rate far exceeding the general adult population.

To draw attention to this issue, the organization used AI to imagine and create what the 27 Club, a group of world-renowned musicians who all died at just twenty-seven years old, might have created if their lives had been longer. The club includes such legends as Jimi Hendrix, Kurt Cobain, Jim Morrison and Amy Winehouse.

Researchers used Google's machine learning platform Magenta to analyze the previously recorded music of each musician in order to form the basis for generating each new track.[3] The AI was fed rhythms, melodies, and lyrics from the artists. This data allowed the AI to produce new songs in the style of the artists themselves, to give a sense of what they might have created had they lived longer.

Once the compositional elements were in place, a production company arranged the musical parts and hired human sound-alike vocalists to record the vocal performances. The Lost Tapes of the 27 Club were distributed publicly on YouTube.[4]

Jukebox AI[5] is a neural network that generates music, including vocals, in a variety of genres and styles of renowned artists. To create the generative model, the JukeBox team crawled the web to curate a dataset of 1.2 million songs (600,000 of which are in English), along with their lyrics and metadata from LyricWiki.[6] The metadata includes the artist, genre, and year of the songs, along with common moods or playlist keywords associated with them. The model can learn over time and generate new music in whatever style the user chooses.[7]

Our research

We designed our quantitative research to mitigate any human bias regarding AI as an artistic creator (not simply used as a tool). Could these AI-generated compositions

perform as well as those generated by their human counterparts when assessed across the four characteristics of artistic value – authenticity, emotion, creativity, and engagement? If AI's musical efforts are to be accepted and embraced by consumers, they will need to go beyond innovative combinations of sounds and motifs.

We partnered with Veritonic, an audio intelligence company, to focus our study on eight specific artists and sixteen musical assets from the selected AI projects.[8] For each artist we studied an original seminal composition that would be known to both casual listeners and fans (control) and an AI-generated composition that emulated their style (test). We collected responses from 1,381 casual listeners and fans in the United States from Aug. 17 to 20, 2022.

Veritonic measures audio assets – advertising creatives, podcasts, musical performances and sonic signatures – and has amassed a large and useful normative database worldwide. The company's ability to measure human-centric attributes at a second by second level of granularity in real time, made it an ideal research partner. Its team mapped our defined characteristics of artistic value with the appropriate set of attributes from Veritonic's platform.

We also composed several follow-up questions to assess how respondents perceived the value of AI-generated music and thus its potential market viability. Unlike the online social postings of the two AI projects we were exploring, which specifically called attention to the music's origin, we deliberately omitted any mention of AI until after the songs had been evaluated.

We recruited discrete groups of roughly 150-200 online respondents, classifying them as either

A summary of the research design and assets follows:

Table 1. Research Design			
Artist	**Assets/Songs (Control)**	**Assets/Songs (AI/Test)**	**AI**
Amy Winehouse	Back to Black	Man I Know	Magenta AI
Kurt Cobain	Come As You Are	Drowned in the Sun	Magenta AI
The Doors	Riders on the Storm	The Roads are Alive	Magenta AI
Jimi Hendrix	Hey Joe	You're Going to Kill Me	Magenta AI
Ella Fitzgerald	Cheek to Cheek	Song in the Style of Ella Fitzgerald	Jukebox AI
Elvis Presley	I Can't Help Falling in Love With You	Song in the Style of Elvis Presley	Jukebox AI
Frank Sinatra	Come Fly With Me	Song in the Style of Frank Sinatra	Jukebox AI
Katy Perry	Last Friday Night	Song in the Style of Katy Perry	Jukebox AI

Table 2. Mapping Characteristics of Artistic Value to Attributes on the Veritonic Platform	
Veritonic Attributes	**Corresponding Characteristic of Artistic Value**
Authentic Trustworthy Familiar	**Authenticity**
Excited Happy Likable	**Emotion**
Innovative Unique	**Creativity**
Average Frequency per Second of All Attribute Interactions	**Engagement**

casual listeners ("I like his/her/ their music but prefer others more") or fans ("one of my favorite artists") of each emulated artist. Each respondent then listened to ninety seconds of audio from each song, original and AI-composed, from the beginning of the first vocal element, rotated to reduce positional bias. Respondents scored the songs at any time, allowing multiple responses, using the following Veritonic attributes:

authentic, excited, familiar, happy, innovative, likable, trustworthy, and unique.

These attributes are captured in every Veritonic analysis and therefore could be compared to the company's normative database of thousands of creative tests that include songs, long form sonic DNAs, music beds, mnemonics and jingles. We then mapped the Veritonic attributes to our characteristics of artistic value. Each value corresponded to two or three attributes.

We presented the songs twice to allow listeners ample opportunity to assess all the attributes. Veritonic calculates and defines the *attribute score* in the following way:

Data is interactively collected while an audio file is being played, measuring the magnitude of selected emotions and descriptors. Each of the Veritonic metrics are on a scale of 0-100.

To measure engagement, we used Veritonic's engagement metric which is calculated differently from its attribute score:

A measure of the frequency of attribute interactions relative to asset duration. This is a metric that counts the number of 'clicks' respondents make on each of the various attributes

Figures 1 and 2: Measuring the Four Characteristics of Artistic Value
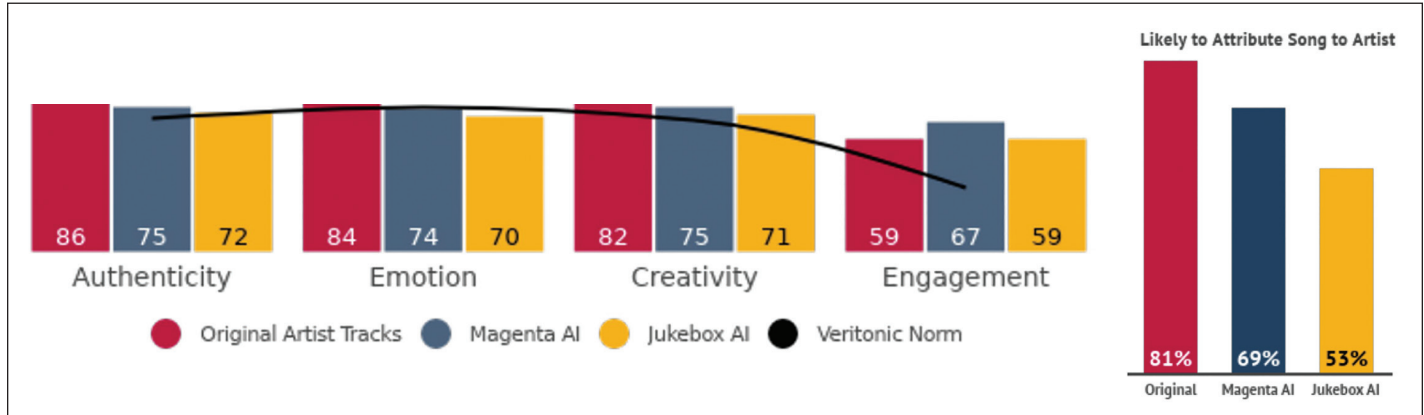
on a second-by-second basis and then expresses a score as an average across all seconds.

For the purposes of this study, we have scaled the engagement scores to Veritonic's attribute scores.

As the audio segment plays, the respondent can select any of the attributes to describe their reactions in real time on a second by second basis. They may also revisit attributes throughout the segment. After they listened, respondents answered a few follow-up questions so that researchers could assess how they perceived the quality of the song, their ability to link the song to the artist, and their intent to further engage with it as a fan, including by streaming or playlisting the track.

At the end of the study, respondents were told that AI had created one of the songs that they had evaluated. They were asked whether they would attribute the AI track to the artist it emulated and rate their level of interest in listening to more AI-generated audio.

What we learned

The AI's compositions performed remarkably well among all respondents across all four characteristics of artistic value. AI-generated offerings scored ten to fourteen points below the artists' original works, but equal to or above the norms for music in Veritonic's global database.

> The AI's compositions performed remarkably well among all respondents across all four characteristics of artistic value.

This is particularly promising given that the AI songs faced a high bar. Its offerings were completely new and unknown to the respondents, while the artist's originals were drawn from their largest commercial successes.

We also observed differences in the respondents' reactions to the works from the two AI generators. The Magenta AI tracks yielded a superior performance across all characteristics. The most striking differences were in the respondents' ability to attribute the tracks to the artists whose work they emulated (69 percent likelihood with Magenta AI vs. 53 percent with Jukebox AI).

Since the Jukebox AI is truly autogenerated, including composition, vocals, and lyrics this result is not surprising, particularly given that its lyrics are occasionally incongruent. Its Frank Sinatra composition, for example, includes the lyrics, *"It's Christmas time and you know what that means… It's hot tub time."* In contrast, Magenta AI's

tracks were composed by the AI, but executed and professionally produced in studios with sound-alike human singers.

While it would be easy to attribute the Magenta advantage to a human element, it is worth noting that even though we observed differences between the two AI-engines, the scores are not appreciably different. Human intervention may have played a role in this study, but it is not unreasonable to believe that future projects or ones with broader scope would yield more comparable results.

Indeed, we observed greater differences between the responses of fans and casual listeners. As Figure 3 shows, fans responded more favorably than casual listeners to both the original tracks and the AI-generated tracks. Across the entire sample, fans made up roughly a third of our respondents, which might account for the lower overall scores when we applied weights.

These results are extremely encouraging for the future of AI-generated music. Self-described fans are likely to be intimately familiar with an artist's style, perspective, and vocal idiosyncrasies. They can detect anomalies or inconsistencies more readily than casual listeners and would be more likely to rebuff imitations.

Figure 3: Fans vs. Casual Listener Response to AI-Generated Music
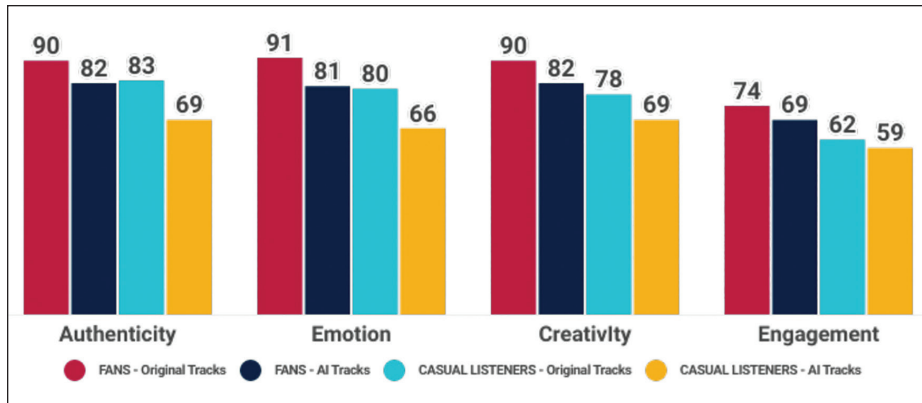
While being exposed to the music during the study, our respondents did not seem to factor these issues into their high scores for authenticity, emotion, or creativity. It was only after they were aware of the AI that fans became less confident of the track's authenticity. Sixty-one percent agreed or strongly agreed that they could not tell the track was generated by AI. We attribute these differences to anti-AI bias.

Toward potential consumer engagement

While the Veritonic platform was able to measure engagement as a compounded interaction metric within a testing platform, true assessments of consumer engagement happen in the open marketplace. To evaluate our last artistic variable, engagement, we questioned consumer intent, a stronger indication of market viability.

In order to get a more realistic response, we asked these questions only after revealing the AI-generation of the songs. After listening to an AI-generated track, 54 percent of respondents indicated that they would like to hear more 'songs like this,' compared to 73 percent for the original tracks. Between a third and just over half of respondents said they would be likely to stream or add the song to a playlist, compared to 67 percent for the original tracks (See Figures 4 and 5).

Qualitative findings and insights

Because both of the AI-generated projects described at the outset have enjoyed widespread distribution on social media platforms, we can understand a lot about listeners' emotional reactions. Where quantitative data allowed us to control for fandom, listenership, and pre-deter-

mined variables, social commentary provided us with a more visceral sense of consumer engagement with the music and the possibilities for AI.

All of the AI-generated tracks of The Lost Tapes of the 27 Club, have been posted to YouTube[9] where the public has unlimited access and can offer qualitative comments for each individual track.

The AI Tracks from JukeBox AI were posted to Soundcloud which also offers unlimited public access for both listening and reposting, again with the opportunity to leave qualitative comments. Table 3 presents our findings for each track by number of streams completed as well as qualitative measures based on 'like' or 'thumbs up' responses and comments. Soundcloud does not offer a 'dislike' or 'thumbs down' option.

Figure 6 charts the relationship of three variables –

1. Consumer Engagement in the form of listens (x-axis),
2. Net Positive Sentiment expressed as likes or dislikes (y-axis), and
3. The raw number of comments (size of the bubble).

Across an array of musical genres from traditional pop to grunge and psychedelic rock, most tracks accumulated massive consumer engagement, from 50,800 to 366,000 in the form of likes and

Figure 4 and 5: Market Viability of AI-Generated Music
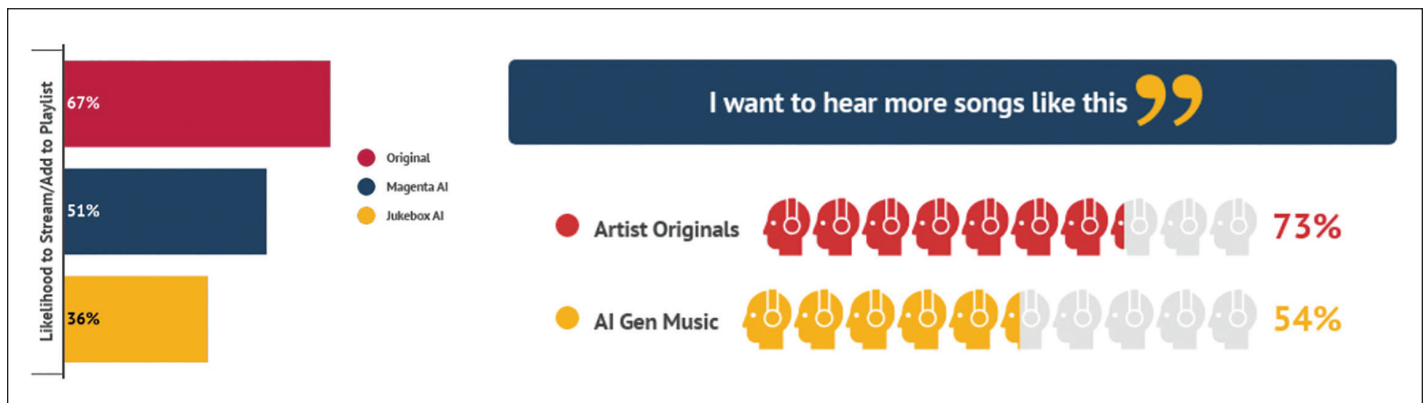
comments. Most listeners found the generative samples to be positive enhancements or interpretations of the original artist's sound. We also found a correlation between high levels of engagement and artists and genres that cross generations such as Frank Sinatra and Elvis Presley as well as pop artist Katy Perry.

Emotion is perhaps the most difficult of our four characteristics to ascribe to AI because it is central to our human identity and admitting that AI can evoke an emotional response challenges our sense of our own value. Nevertheless, some comments revealed that the AI did evoke emotional responses:

"Damn… We are now to the point where AI can make me feel something…"

Figure 7 offers a sampling of comments, organized by our four

Table 3. Summary of Qualitative Data Sources						
AI Platform	Track Artist	Track Name	Listens	Likes	Dislikes	Comments
Magenta AI	Jimi Hendrix	You're Going to Kill Me	68,329	1,100	16	175
Magenta AI	Amy Winehouse	Man I Know	146,203	3,400	37	415
Magenta AI	The Doors	The Roads are Alive	50, 817	84	45	198
Magenta AI	Nirvana	Drowned in the Sun	87,222	424	103	208
JukeBox AI	Elvis Presley	Rock	251,000	257	n/a	62
JukeBox AI	Frank Sinatra	Pop	366,000	1060	n/a	206
JukeBox AI	Katy Perry	Pop	256,000	483	n/a	89

characteristics, with the source artist noted in parentheses. The responses demonstrate a healthy mix of awe and skepticism.

In light of these findings and the immense library and lexicon of music which has been digitized over the past several decades, the groundwork has been laid for new AI musicians to emerge.

As capabilities become more sophisticated, the commercial opportunities for consumer engagement in the music industry could expand exponentially as music discovery is no longer relegated to new or previously unknown artists, but to expanding the catalog of historical artists as well. Of course, the applications,

Figure 6: Social Media Response to AI-Generated Music
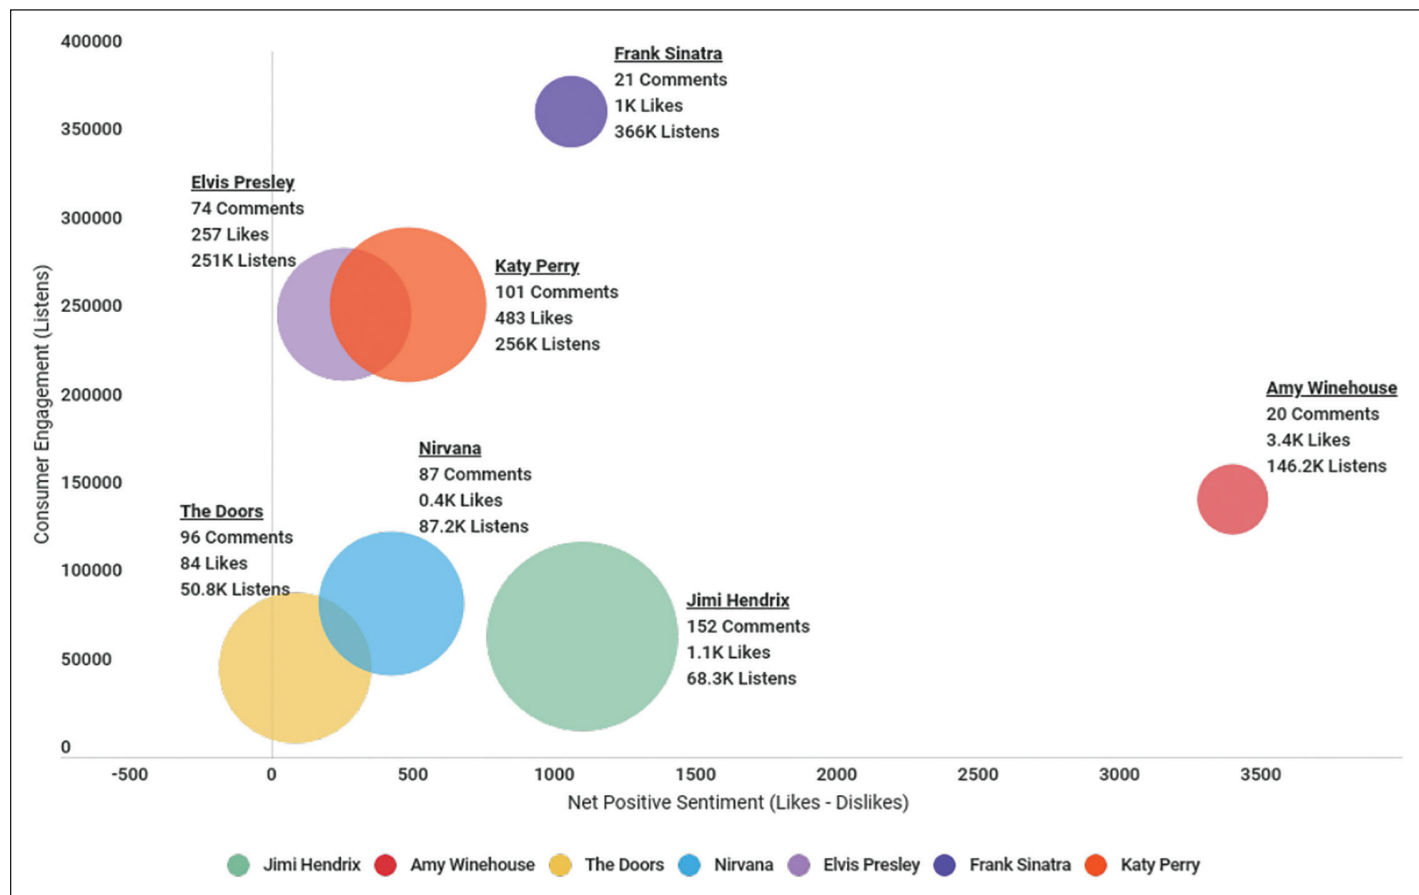
EMOTION

This genuinely made me tear up. (Katy Perry)

How the hell does this have so much emotion? (Nirvana)

I come back to listen to this like I'm mourning a loved one. (Amy Winehouse)

Wow... for the first time it all makes sense how an AI could possibly evolve to have emotion (Amy Winehouse)

A song's never made me cry so much (Nirvana)

OMG why is this making me a little emotional? (Katy Perry)

How is this song so deep in meaning and emotion? I'm terrified. (Nirvana)

Where is the soul? It is clever but you can tell that it is not quite right. (Sinatra)

AUTHENTICITY

This is like a window into a parallel universe where Kurt never died and it make me wonder the beautiful songs we would've had if he was still alive (Nirvana)

It's like AI went to heaven and said "Yo Kurt, you got anything new?" (Nirvana)

I'm super frightened and pumped by this ... this song goes too hard to be a soulless computer (Jimi Hendrix)

I didn't know Skynet was doing covers (Nirvana)

That's how a robot would perceive art. (Katy Perry)

It's jibberish. It sounds like her but none of the soul. (Amy Winehouse)

I'm actually really glad to find out AI still has no true grasp on what music is really about. (Sinatra)

Don't give up your day job (Elvis)

Gosh this is rubbish. Frank sang music not muzak. (Sinatra)

This is insane. I had it in my head that AI wouldn't be able to write compelling music for at least another ten years.
Guess I was wrong (Nirvana)

If AI makes music like this they can rule the world (Amy Winehouse)

OMG!! The power chords and progressions are amazing! (Katy Perry)

The intro is an emotion on its own. (Nirvana)

2:28 ... I want that melody. (Nirvana)

If the AI came up with the lyrics on its own I'd say it's time to craft it a body and give it citizenship! (Elvis)

We need a human to write lyrics. (Amy Winehouse)

What do you call a song that refuses to exist in a key? (Sinatra)

We're in the early days of synthetic media. (Elvis)

CREATIVITY

Where can I buy the album (Katy Perry)

New Katy Perry single confirmed (Katy Perry)

I heard this on Joe Rogan's podcast and had to find the whole song. This is amazing. (Amy Winehouse)

I hope they release an album of AI songs generated by AI. (Nirvana)

It would be so cool to see a completely new album for Nirvana with songs just like these. (Nirvana)

This makes me happy as a consumer. It just means the unmade music we will never hear will be heard and it just gives us more noise to love (Amy Winehouse)

Soundcloud REALLY needs a dislike button (Sinatra)

ENGAGEMENT

ethics, and questions that these capabilities raise are equally numerous.

Looking ahead: The culturally curious will lead the way

These two AI platforms, when set to learn and reproduce the works and stylistic nuances of legendary musicians such as Jimi Hendrix, Kurt Cobain, and Jim Morrison generated enormous consumer engagement through social media.

Today most listeners view The Lost Tapes of the 27 and other AI-generated music experiments as novelties. However, as AI/ML continues to grow and learn from an ever increasing archive of music, these innovative applications demonstrate that we are on the cusp of what will be possible as AI is more extensively applied to music, providing entertainment and attracting customer engagement.

The quantitative data results suggest an openness to new AI-generated works rather than what might have been feared by self-appointed afficionados or other cultural gatekeepers. Certainly, the comparative data from our four characteristics of artistic value demonstrate an acceptable level of accomplishment to upwards of two-thirds of music fans.

There can be no doubt that we will hear ever more refined renditions of what legendary composers might have produced had they lived longer. And from a business perspective, a wealth of new compositions wait to be rendered and monetized as AI/ML is applied to music for every purpose, from pure consumer entertainment to advertising, television, motion pictures, and other commercial enterprises. Finer slices of addressability across all media can be expected to increase demand for creative assets beyond human capacity, necessitating further evolution of creative iteration on demand.

> There can be no doubt that we will hear ever more refined renditions of what legendary composers might have produced had they lived longer.

Within the music industry specifically, we can already see both artist-to-consumer and business-to-business applications for this technology. AI-generated works that can both imitate and iterate the compositions of human artists without copyright infringement will largely alleviate the cost of creative rights for business applications. Meanwhile, as copyright royalties dwindle, artists will look to AI to help them increase their earnings.

AI presents artists with an enormous opportunity to lean into the coming uniquity movement, in which artistic works can be created in pieces and then assembled in infinite combinations, producing bespoke masterworks for a premium price. For music fans, libraries and playlists are about to expand exponentially as 'New Music Monday' may include everything from the latest Billie Eilish single to the newest AI evolution of the Beatles catalog.

Finally, AI will not be working alone in the next iteration of creativity, especially as humans, avatars, and corporations are already intermingling in the metaverse. Creative humans are already exploring innovative applications that surprise and delight the culturally curious. Expect musicians and visual artists to stretch boundaries and adopt new AI partners as they strive to express the duality of life lived both actually and virtually.

Our findings suggest that consumer engagement with AI-generated music is all but certain to increase. We confidently expect the further refinement of this technology to expand this engagement, generating opportunities for both artists and marketers. For now, with the help of AI/ML, the legendary composers, artists, and musicians explored in our study may be gone, but their music, both actual and generative, will definitely not be forgotten.[10] ■

Author Bios

Alan Schulman is cofounder and managing partner of UpperRight. As a chief creative officer in marketing and a trained jazz musician, he is a sought-after advisor in applying music to brand ethos, content, and personality. He has received numerous awards in advertising and design and holds degrees in journalism and communications from Ohio State University and Northwestern University, and an MFA from Howard University. He is a voting member of the Recording Academy (Grammys).

Stacey Lynn Schulman is the founder of Hi: Human Insight which focuses on human aspects of data and strategy. Her career spans research leadership at global media and advertising agencies. She is routinely quoted as a recognized fan culture expert and was the first researcher inducted into the American Advertising Federation Hall of Achievement. She maintains a career as a Top 10 Billboard jazz vocalist and is a voting member of the Recording Academy (Grammys).

Endnotes

1. Charley, Pease and Colton, 2012
2. https://losttapesofthe27club.com/
3. https://magenta.tensorflow.org/
4. https://losttapesofthe27club.com/
5. https://openai.com/research/jukebox
6. https://lyrics.fandom.com/wiki/LyricWiki
7. For a more comprehensive paper and breakdown of JukeBox AI and the many styles and artists it has been applied to, visit https://openai.com/research/jukebox Jukebox AI compositions can be heard on the Soundcloud® digital distribution platform here.
8. https://www.veritonic.com/
9. https://www.youtube.com/watch?v=zPkOtW5n8-E
10. Detailed results of this study may be found at https://www.hihumaninsight.com/aimusicstudy

Wizenoze: The Value of Engaging the Customer's Customer

Stefano Puntoni, The Wharton School, University of Pennsylvania

Stefano Puntoni examines how Wizenoze, a Dutch educational technology startup, uses AI to match educational content to each learner's reading skills. The case illustrates how AI can allow service personalization at scale in order to increase user engagement and satisfaction.

Customer engagement is a priority for many companies, but what it looks like in practice depends on the product or service. For educational technology (ed tech) companies, customer engagement requires that users pay attention to the product or service in order to benefit from it. We use technology to facilitate learning and to ensure that students are engaged and ready to process and retain information. Artificial intelligence (AI) can help us to achieve this goal deeply and at scale.

Consider the experience of Amsterdam-based start-up Wizenoze. It was founded in 2013 by Diane Janknegt and Theo Huibers whose purpose was to

curate information on the Internet for children and young students to ensure that what they see is age appropriate, true, safe, and appropriate to their educational reading level.[1] A research grant from the European Union to academics at the University of Twente in the Netherlands, where Huibers is a computer science professor, got the start-up off the ground.

Janknegt and Huibers believed the Internet offers an incredible wealth of data but that finding relevant, reliable, and readable information for students from primary to graduate school is difficult. They knew that search engines like Google cannot match content to the reading skills of the user and do not always yield results suitable for quality education. They founded Wizenoze to offer a better way.

The company developed algorithms that matched online content, both free and proprietary, to a student's learning goals and reading skills. Humans are in the loop to screen the information for relevance and readability. This technology is integrated into a school's learning system and matched to the curriculum. Teachers and students can then search the Wizenoze database for reliable and relevant data as they teach or learn.

Engaging the Customer's Customer

But in 2019, revenues were not meeting expectations and the company decided to radically change direction. Janknegt, the CEO, describes this pivotal moment as "running into a wall while driving 100 kilometers per hour." She realized they had made a major mistake. In their product development, the company's leaders had listened carefully to their business to business (B2B) customers — learning management system providers (such as Canvas, Moodle) and publishers (Pearson, Encyclopedia Britannica, and the like).

But they had not listened to end users: students and teachers. What they had failed to realize was that even if Wizenoze gave students access to a smaller, curated trove of information than the open Internet, many students still found searching for specific information too difficult. They wanted an easier way to get the best information about their topic.

Wizenoze quickly changed its proposition. Instead of being a search engine for education, the company dropped the search bar altogether. Instead it created a tree of curated libraries of educational content. A student interested in learning about a particular topic in biology could now click on 'biology' to go immediately to a high-quality library of relevant information selected by Wizenoze to match their reading skills. From within the biology library, the student could then click 'cell development' to refine the results.

AI Application

The approach taken by Wizenoze differs from that used by most search engines, which judge content eligibility based on blacklisting. In this approach, users can access most sites except those that are blacklisted. Wizenoze instead whitelists websites, blocking websites by default unless their content is proven to be trustworthy and suitable.

Humans determine the trustworthiness of each source, assuming a trusted source to have trustworthy information. Meanwhile algorithms match materials to the required reading level and curriculum. Humans provide additional oversight by checking the content libraries.

How Wizenoze curates and checks the content:

1) Web crawlers browse the Internet to identify content potentially suitable for education.

2) Human experts, such as teachers, select trusted sources from the candidate pool of websites.

3) A combination of human and automated agents strip the selected websites of content that is not deemed educational.

4) Machine learning algorithms assign a reading level to the content and then match it to the curriculum.

To match web content from its collection with learning objectives in the curriculum, Wizenoze fine-tuned a state-of-the-art neural information retrieval (NIR) model using a Siamese BERT network.

Each curriculum in the system is defined by a structure that is an ordered sequence of layers. Different curricula might have different structures composed of layers such as 'Grade,' 'Subject,' 'Unit,' 'Topic,' and 'Query.' These algorithms are a crucial source of competitive advantage for Wizenoze.

5) The top results are then checked by human experts before being included in the curated content library.

6) A web crawler ensures that the information is kept up-to-date.

Customer Engagement Results

Much online information is not suitable for education because it is untrustworthy, irrelevant, or difficult to understand. Furthermore, it is hard for many students to formulate effective queries and judge which websites to trust. The Internet's enormous potential to support and facilitate education is therefore still largely unfulfilled.

Wizenoze provides a solution to this important problem. With Wizenoze, students can log into their school's online learning site and find the best online resources on any subject in the curriculum. The system will then curate this content according to relevance and reading level.

To gauge its system's effectiveness, Wizenoze conducted a study in the United Kingdom (UK) to measure customer engagement.[2] It tested the tool on a focus group of teachers and 190 students at different reading levels. The teachers reported benefitting from the tool's support of students in a variety of tasks, including reviews (recalling topics previously discussed in class), research (for presentations or reports), exploring topics related to those covered in class, and helping students individually with specific subjects.

Students also responded positively to this new learning tool while their teachers appreciated the ability to personalize, for example, by choosing how many results appeared on a page. For a UK client, the company also conducted a later study using four English primary schools in Sussex, London, and Bristol. This unpublished study confirmed the effectiveness of Wizenoze content libraries for customer engagement and learning.

In the latter study, 87 percent of teachers said the search results improved for their pupils in terms of relevance, readability, reliability, ease of use, and enjoyment. Students confirmed the teachers' assessments, rating Wizenoze as better than Google for finding relevant information (average of 4.8 on a 5-point scale). Although long-term data on the service is not available at this time, students indicated an interest in using it if it became available in their school (average of 4.1 on a 5-point scale).

The company is currently growing rapidly and already reaches over a million students. It is expanding internationally, and expects its revenues to double annually for the next three years, with a growing global footprint that includes the UK, Europe, the Middle East, and India.

Lessons to Take Home

The case study of Wizenoze, and particularly its pivot from open searches to curated libraries in 2019, offers important lessons for AI-driven companies seeking to boost customer engagement:

1. *Value chain perspective on B2B marketing.* Wizenoze had interacted frequently with its B2B clients and tried to be as responsive as possible to their needs. However, this customer focus was not sufficient to ignite growth. The company's trajectory only picked up once it started paying attention to end users. B2B companies must take a value chain approach to product design and value generation. In short, caring about your customers is not enough; your customers' customers are equally important.

2. *The value of information in a post-truth world.* Even in a world where search costs are essentially zero, there is room for businesses to grow by helping people find information online. The Wizenoze case suggests a productive way to think about new opportunities in digital business: Look for needs that are currently unmet by ad-driven business models. We glorify big data, but small good data beats big bad data. Wizenoze's differentiation is a shift from information quantity to information quality.

3. *Customer engagement as an antecedent to providing value.*

In many industries, customer engagement is directly linked to benefits. In education, it is crucial that students have access to materials that match their reading skills and learning goals. Those who do not find it difficult to process the information provided and struggle to stay motivated. AI can facilitate engagement in education because it enables personalization at scale.

4. *Boosting engagement by reducing choice.* Behavioral science has long established that people value having a lot of choices. But in fact, they can also suffer from having too many choices. Talking to end users helped the Wizenoze team realize that, for many students, searching for information and sorting through search results is cumbersome. Wizenoze became more successful by reducing the number of choices. Leaders should always remember that simplifying choices can improve the customer experience and increase customer engagement. Sometimes, less is more. ◼

Author Bio

Stefano Puntoni is the Sebastian S. Kresge Professor of Marketing at The Wharton School, University of Pennsylvania. His research investigates how new technology is changing consumption and society. Puntoni teaches marketing strategy, new technologies, brand management, and decision making.

Endnotes

1. The article is based on interviews and research conducted for "Wizenoze: Matching Digital Content to Learners in a Post-truth World" by Kassiani Nikolopoulou and Stefano Puntoni, published by the Case Development Centre at the Rotterdam School of Management, Erasmus University. https://www.rsm.nl/cdc/multimedia-cases/wizenoze/

2. Jochmann-Mannak, H., C. Frasera, G. de Vries, T. Westerveld, and T. Huibers (2021), "Supporting self-directed learning and exploratory research on the web," European Conference on Technology Enhanced Learning for a Free, Safe, and Sustainable World, Bolzano, Italy.

Customers who trust products and services engage more, which improves their experience and satisfaction. AI tools now shape many customer decisions, so their trustworthiness is critically important. Jennifer Shkabatur and Alex Mintz explore how commonly accepted measures of trustworthiness in AI can be practically tested and ranked.

Developing a Trustworthy AI Rating System and Its Impact on Customer Engagement

Jennifer Shkabatur
Reichman University

Alex Mintz
Reichman University

Every day artificial intelligence (AI) products shape a broad array of customer decisions about finance, e-commerce, health care, leisure, professional recommendations, and more. Indeed, many businesses have drastically changed their modes of operation, deploying AI so they can better understand their customer's preferences, shape their decisions and behavior, and strengthen their engagement.[1]

Since customers' trust in products drives their engagement, the question of whether AI products can be trusted is critically important.

Since customers' trust in products drives their engagement, the question of whether AI products can

be trusted is critically important. We have devised a method by which commonly accepted measures of trustworthy AI can be practically tested, and a ranking scale for trustworthiness in various sectors. This method can be applied by the business community, investors, regulators, ranking agencies, and the customers themselves.

Trustworthiness as a driver of customer engagement

A longstanding and primary task of marketers is to persuade consumers who are on the fence about a product or service to go ahead and buy.[2] In order to achieve this goal, they must build trust — confidence in the product's reliability, robustness, security, integrity, and ability to meet their needs.[3]

Trust is therefore generally recognized as one of the primary factors that strengthen customer engagement, sales, and satisfaction.[4]

Trust contributes to a customer's willingness to purchase the product again, builds a deeper and more personal connection with the product, and thus increases sales[5] and, ultimately, the probability of future repurchases.[6] Conversely, lack of trust in a product reduces engagement in both individual and business customers.[7]

To increase customer engagement in all interactions, businesses use a range of AI products, including recommendation systems, conversational agents, sentiment analyzers, and natural language processing algorithms. While customers draw huge benefits from these AI products, they may also find assessing their trustworthiness to be challenging for several reasons.

1. The opacity of AI operations tends to prevent public scrutiny, particularly since AI systems are difficult to explain and understand.[8] Customers may respond to this uncertainty with bias and discrimination that

Table 1. Assessment Questions for Trustworthy AI	
1. Human agency and oversight	• *Self-assessment.* Does the AI product allow the user to reasonably assess or challenge the product? • *Intervention.* Does the AI product enable human intervention during its decision cycles? • *Discretion.* Is it possible to integrate human discretion into the operation of the AI product? • *Override.* Is it possible for a human to override a decision made by the product?
2. Technical robustness and safety	• *Effectiveness.* Does the AI product achieve the outcomes that it promises to achieve? • *Accuracy.* Does the AI product make correct predictions, recommendations, or decisions? • *Safety.* Is the AI product safe to use (in its data security protocols, and software and hardware safety)? • *Reproducibility.* Does the AI product exhibit the same behavior when applied under similar conditions?
3. Privacy, data governance, and legal compliance	• *Data sources.* What are the data sources used by the AI product? How and when was this data collected? Was consent acquired? How is this data used by the AI product? How is the integrity of the data ensured? • *Data access and management.* Who has access to the data utilized by the AI product, and under what conditions? Who is eligible to change or manage the data? • *Legal compliance.* Does the AI product comply with the legal requirements of the country in which it is deployed (e.g., privacy, health and safety, etc.)?
4. Transparency and explicability	• *Traceability.* Are the data and processes that yield the AI product's decision documented? Is it possible to trace back and link the product's inputs and outputs? • *Explicability.* Is it possible to explain in plain language how the AI product works? • *Communication.* Is any information provided regarding the limitations of the AI product, its level of accuracy, or other issues related to its operation?
5. Diversity, non-discrimination, fairness	• *Biases.* Can the data used by the AI product be affected by inadvertent bias, discriminatory patterns against certain groups, or incompleteness? Can the outcomes of the AI product lead to discrimination against certain groups? • *Accessibility.* Is the AI product user-centric and designed in a way that allows anyone to use it, regardless of age, gender, abilities, or characteristics?
6. Societal and environmental well-being	• *Sustainability.* What are the environmental impacts of the AI product (including its development, deployment, and usage)? What are the resource usage and energy consumption during the product's training and operation? • *Social impacts.* Can the AI product adversely affect users' mental or social well-being?
7. Accountability	• *Auditability.* How practical is it to audit the AI product? Can it be independently audited? • *Minimizing and reporting negative impacts.* How can the AI product's actions or decisions be reported? • *Redress.* If the AI product causes adverse effects, what redress venues are available?

Figure 1. Assessing Three AI Products

We analyzed and rated for trustworthiness the following products:[15]

1) **Fitness coaching app** that offers personalized coaching for improving fitness and mindfulness, as well as developing healthy lifestyle habits. Users indicate the goals they want to achieve and answer a survey about their habits. With this information, the app generates a personalized training plan for each user.

2) **AI-backed search engine** that analyzes various types of financial documents, relies on natural language processing for improved results, and transforms unstructured data into structured. It then allows users to search for specific terms within the documents based on geography and time. It includes features such as 'synonyms search,' and sentiment analysis, and prioritizes results according to importance and relevance.

3) **Conversational AI-based virtual assistant** that specializes in banking. It helps users analyze their financial activities, generating information and recommendations based on various factors and prior interactions.

companies struggle to reverse or even uncover.

2. Low public awareness and AI literacy, which may lead the public to neglect and disregard the risks associated with AI.[9] A recent survey showed that only 62 percent of respondents have any knowledge of AI and the majority reported that they had low understanding of AI. When presented with a range of common AI applications, respondents were not aware that the described technology used AI.[10]

3. Rapid scaled deployment of AI that hinders companies in properly scrutinizing AI products before investing in and deploying them.[11]

4. Lagging regulation of AI systems such that governments cannot provide timely and adequate legal recourse when AI problems arise, complicated by differences in the rules of separate countries.[12]

A vast range of stakeholders, including private companies, academics, government agencies, intergovernmental organizations, and professional associations, have formulated principles for ethical, trustworthy, and human-centered AI systems that customers could trust. By examining these principles, we traced significant commonalities and emerging norms.

In 2019, the European Commission (EC) proposed a seven-principle framework for 'trustworthy AI' systems that has been broadly accepted:[13]

1. **Human agency and oversight:** AI systems should empower human beings, allowing them to make informed decisions and fostering their fundamental rights.

2. **Technical robustness and safety:** AI systems need to be resilient and secure, offering their users accurate, reliable, and effective services.

3. **Privacy and data governance:** Full respect for privacy and data protection should be ensured.

4. **Transparency and explicability:** AI systems and their decisions should be explained in a manner adapted to the stakeholder concerned.

5. **Diversity, non-discrimination, and fairness:** Unfair bias and outputs that discriminate against specific groups must be avoided.

6. **Societal and environmental well-being:** AI systems should benefit all human beings, including future generations. They must be environmentally friendly and sustainable.

7. **Accountability:** Mechanisms should be put in place to ensure responsibility and accountability for AI systems and their outcomes.

Many subsequent AI ethics frameworks focus and draw upon these principles.[14] Still, while we have achieved notable unity as to the characteristics of trustworthy AI products, few governments or organizations use those characteristics to actually assess AI products. We have therefore developed a procedure for testing AI products against the EC's parameters.

Assessing the trustworthiness of AI products

We begin by translating these parameters into questions which we then use to assess three real AI products in common areas of customer engagement: finance, health and fitness, and natural language processing. The resulting ratings offer a standardized assessment of trustworthiness across AI products which can be used by consumers, regulators, businesses, and investors and which therefore can be a significant driver of customer engagement.

We measure the trustworthiness of AI products based on the questions in Table 1. These questions reflect the EC's definitions and also closely follow the wording used in other studies of trustworthy or ethical AI.

This straightforward scale offers businesses and consumers an easy, standardized, and convenient way to assess the trustworthiness of an AI product.

Table 2. Trustworthiness Assessment and Rating

	Personalized coaching app	AI-backed search engine	Conversational AI-based virtual assistant
Human agency and oversight	The coaching plan's objectives are determined by the user. Changing one's preferences results in changed outputs.	The outputs are modified interactively in response to user changes. But the user cannot challenge the system or establish discretion as part of its operation.	The chatbot is interactive. If the bot cannot understand the user's inputs, it suggests chatting with a live agent.
Technical robustness and safety	The product reflects the user's preferences, yet it does not claim or prove that its recommendations are optimal. It exhibits similar behavior when repeated under similar conditions.	The system displays the data that it finds, so accuracy is easy to establish. However, there may be missing results, and no convenient way to account for these.	The system does not undertake specific actions but rather provides recommendations. There is no information on whether these are optimal.
Privacy, data governance, and legal compliance	The privacy policy is very clear, complying with the laws of California. It states what, how, and when data is collected; requires user's consent for use of the data; and explains under what conditions data can be shared with third parties.	The system's big data is not related to its users. It accesses extremely limited user information (e.g., email and name), and is highly secured.	The privacy policy elaborates on user data usage. However, while user data is used to improve the service, it is not explained how data integrity is maintained.
Transparency and explicability	The system's outputs are explained, but does not explain why these outputs were preferred over others. There is no information regarding limitations, and no way to trace back to its operations mode.	The system highlights the keywords and search terms that it found, yet there is no explanation of why and how certain search terms are prioritized. The sentiment analysis is also not explained.	The bot responds as part of a conversation, but there is no explanation of how financial recommendations are made.
Diversity, non-discrimination, fairness	The outputs are based on user-specified goals and body characteristics. However, it is not clear which biases may emerge against certain user groups.	Some AI features do not support non-English content. The system focuses on financial documents and there seems to be no discrimination against certain user groups, yet there is no clear way to test this.	No sufficient information for a response. Some user groups may be discriminated against by the virtual assistant. There is no way to compare financial recommendations provided to different users.
Societal and environmental well-being	No information regarding environmental impacts. Since it is a fitness app, it is expected to positively affect users' well-being.	The system uses Amazon Data Centers. Its general societal impact seems to be neutral.	No information, but the system does not appear to raise considerable concerns.
Accountability	Users can change the suggested workout plan. They can also contact the company's support desk. It is unclear whether and how redress for adverse outcomes could be provided.	There is no practical way to audit the system, but users can contact the company's support desk. It is unclear whether and how redress for adverse outcomes could be provided.	There is no practical way to audit, but users can contact the company's support desk. It is unclear whether and how redress for adverse outcomes could be provided.

We used these questions to assess the trustworthiness of real AI products and developed a trustworthy AI scale ranging from green (full compliance with the principles), to yellow (partial compliance), to red (lack of compliance and/or impossible to check). This straightforward scale offers businesses and consumers an easy, standardized, and convenient way to assess the trustworthiness of an AI product and, as a result, its impact on customer engagement.

We pilot tested the proposed method on three real AI products, which were randomly selected from a database of 1,700 that we assembled from key sectors that use AI. Dr. Shkabatur and an independent, highly skilled computer scientist conducted the test by examining the products' websites and running demos. Using the questions in Table 1, we rated the products as green, yellow, or red.

Figure 1 provides general information about the AI systems. Their trustworthiness ratings appear in Table 2.

This pilot examination of our rating method gives a sense of whether the trustworthy AI parameters could be applied to real AI products. It has led us to the following reflections:

1) Much debate has been dedicated to privacy and data governance concerns about AI products. Our analysis suggests that the designers of these products have a considerable awareness of privacy requirements and expectations, and that they could feasibly test for and achieve compliance with privacy protection guidelines. Assessing technical robustness also does not appear, by this initial analysis, to present particular difficulties for designers of AI products that target general customer markets.

2) It is possible to examine whether the algorithms in AI products are discriminatory and produce biased results.

3) We found it tougher to assess trustworthiness principles that require peering into the AI products' black box. Transparency and explicability are more difficult for products that operate on more complex datasets. Accountability presents a challenge, because AI products typically do not provide information on whether they could be externally audited, or how negative impacts could be redressed. This generates a particular challenge for customer engagement since AI systems do not typically provide sufficient information on whether customers can trust them to achieve fair and transparent outcomes.

4) Assessing the societal and environmental well-being of these AI products is not easy, raising hurdles for ethical customer engagement. Generally, neither consumers nor businesses know or have the tools to assess how the AI system affects users. Measuring those affects would again require opening the black box.

A final word

Given the widespread use of AI products for customer engagement and the growing reliance on them among businesses and customers, the trustworthiness of these products must be thoroughly scrutinized by investors, businesses, regula-tors, and the users themselves. Indeed, the question of whether AI products can be trusted is critical and has received considerable attention.[16]

Our procedure for rating the trustworthiness of AI products is rooted in a practical application of commonly accepted EC parameters. It measures the trustworthiness of AI products on a straightforward scale of green-yellow-red, which could be readily understood by customers. This scale could work in a way similar to the hygiene ratings that restaurants display in their front windows. An impartial authority such as a ranking agency, academic body, or the AI company itself could administer the ranking process, provided it offered a full explanation on each parameter.

While an external examination can effectively assess certain common parameters of trustworthy AI, others require opening the black box, an endeavor that will require further consideration. We suggest that only products graded green across all parameters should be considered fully trustworthy, while those with some yellow grades should be viewed as partially trustworthy.

> Our procedure for rating the trustworthiness of AI products is rooted in a practical application of commonly accepted EC parameters.

Author Bios

Jennifer Shkabatur is an Assistant Professor at the Lauder School of Government, Diplomacy & Strategy and Senior Researcher at the Computerized Decision-Making Lab at Reichman University in Israel. Her research focuses on the ethics of artificial intelligence, innovation in the digital economy, and data and information policies. She also consults on these issues with the World Bank, UNDP, ILO, and more. She earned doctorate and master's degrees from Harvard Law School. jshkabatur@runi.ac.il

Alex Mintz, is Director of the Computerized Decision Making Lab at Reichman University. He was awarded the Lifetime Achievement Award from the Israeli Political Science Association in 2019 and has received numerous other accolades. Mintz was editor in chief of the journal *Political Psychology* and on numerous editorial boards. Among his many books are: *The Polythink Syndrome: U.S. Foreign Policy Decisions on 9/11, Afghanistan, Iraq, Syria, Iran,* and *ISIS,* and *Understanding Foreign Policy Decision Making.*

Endnotes

1. See:
   Kumar, V., Rajan, B., Venkatesan, R., & Lecinski, J. (2019). "Understanding the role of artificial intelligence in personalized engagement marketing," *California Management Review*, 61(4):135–155.
   Davenport T., Guha A., Grewal D., Bressgott T. (2020). "How artificial intelligence will change the future of marketing," Journal of the Academy of Marketing Science, 48 (1):24-42.
   Bag, S., Srivastava, G., Al Bashir, M. M., Kumari, S., Giannakis, M., & Chowdhury, A. H. (2021). "Journey of customers in this digital era: Understanding the role of artificial intelligence technologies in user engagement and conversion." *Benchmarking*.

2. See:
   Morgan, R.M., and Hunt, S.D. (1994). "The commitment-trust theory of relationship marketing," *Journal of Marketing* 58:20–38.
   Sichtmann, C. (2007). "An analysis of antecedents and consequences of trust in a corporate brand," *European Journal of Marketing* 41(9/10):999–1015.

3. Delgado-Ballester, E., Munuera-Aleman J.L., Yague-Guillen, M.J. (2003). "Development and validation of a trust scale," *Int'l J. Market Research* 45(1):35-58.

4. See:
   Choung, H., David, P., & Ross, A. (2022). "Trust in AI and its role in the acceptance of AI technologies." *Int'l Jour. Human–Computer Interaction*,1–13.
   Ferrario, A., Loi, M., Vigano, E. (2019). "AI we trust incrementally: A multi-layer model of trust to analyze human-artificial intelligence interactions," *Philosophy & Technology*, 1-17.

Wang, X., Tajvidi, M., Lin, X., Hajli, N. (2019). "Towards an ethical and trustworthy social commerce community for brand value co-creation: A trust-commitment perspective." *J. Bus. Ethics*, 1-16.
Shin, D. (2020). "User Perceptions of Algorithmic Decisions in the Personalized AI System: Perceptual Evaluation of Fairness, Accountability, Transparency, and Explainability," *Journal of Broadcasting & Electronic Media*, 64(4):541-565.

5. See:
   Bowden, J. L-H. (2009). "The Process of Customer Engagement: A Conceptual Framework," Journal of Marketing Theory and Practice,17:1, 63-74.
   Singh, J.J., Iglesias, O. & Batista-Foguet, J.M. (2012). "Does Having an Ethical Brand Matter? The Influence of Consumer Perceived Ethicality on Trust, Affect and Loyalty," *J. Bus. Ethics*. 111:541–549.

6. Madhani, P. M. (2020), "Ethics in Sales and Marketing: Key Advantages", *Marketing Mastermind*, 17(5):53-58.

7. See:
   Kingshott, R.P., Sharma, P., Chung, H.F. (2018). "The impact of relational versus technological resources on e-loyalty," *Industrial Marketing Management*, 72:48-58.
   Grewal, D., Guha, A., Satornino, C.B. and Schweiger, E.B. (2021), "Artificial intelligence: the light and the darkness," *J. Bus. Research*, 136:229-236.

8. See:
   Rudin, C. (2019), "Stop explaining black-box machine learning models for high-stakes decisions and use interpretable models instead," *Nat. Mach. Intell.* 1:206–215.
   Rai, A. (2020), "Explainable AI: from black box to glass box," *J. Acad. Mark. Sci.* 48:137–141.

Guha, A., Grewal, D., et al. (2021). "How artificial intelligence will affect the future of retailing." *Journal of Retailing*, 97(1):28-41

9. Curtis, C., Gillespie, N. & Lockey, S. (2022). AI-deploying organizations are key to addressing 'perfect storm' of AI risks. *AI Ethics*.

10. Gillespie, N., Lockey, S., Curtis, C. (2021). Trust in Artificial Intelligence: a five-country study. *The University of Queensland and KPMG Australia*.

11. Zhang, D., Mishra, S., Brynjolfsson, E., Etchemendy, J., Ganguli, D., Grosz, B., Lyons, T., Manyika, J., Niebles, J. C., Sellitto, M., Shoham, Y., Clark, J., Perrault, R. (2021). "*The AI Index 2021 Annual Report*".

12. See:
   Mittelstadt, B. (2019), "Principles alone cannot guarantee ethical AI," *Nat. Mach. Intell.* 1:501–507.
   Taeihagh, A. (2021). "Governance of artificial intelligence." *Policy and Society*, 40(2):137-157.
   Djeffal, C., Siewert M. B., Wurster S. (2022). "Role of the state and responsibility in governing artificial intelligence: a comparative analysis of AI strategies," *Journal of European Public Policy*.

13. These principles rely on four core ethical principles that lie in the basis of the EU Charter: respect for human autonomy; prevention of harm; fairness; and explicability.
   European Commission (2019). "Ethics guidelines for trustworthy AI." *EC HLEG* https://digital-strategy-ec-europa-eu.ezprimo1.idc.ac.il/en/library/ethics-guidelines-trustworthy-ai.

14. The OECD developed a typology, which includes roughly similar principles ("Inclusive growth, sustainable development and well-being;" "Human–centred values and fairness;" "Transparency and explainability;" "Robustness, security and safety;" and "Accountability.").
   OECD/LEGAL/0449. (2019). "Recommendation of the Council on Artificial Intelligence." OECD.
   The IEEE's Ethically Aligned Design for Autonomous and Intelligent Systems resulted in comparable principles.
   IEEE (2019) "Ethically Aligned Design – First Edition", *The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems*.
   See also: Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., & Srikumar, M. (2020). *Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches.*, which compares 36 distinct initiatives to define "ethical" and "rights-based" AI and identifying largely similar themes.

15. The products were tested in July 2022.

16. See:
   Ameen, N., Tarhini, A., Reppel, A., & Anand, A. (2021). "Customer experiences in the age of artificial intelligence." *Comp. Human Behavior*, 114:106548.
   Dwivedi, Y.K, Hughes L., et al. (2019). "Artificial Intelligence: Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy." *Int'l Jour. Info. Management*.

Book Review: Seven Ways that AI Will Transform Customer Engagement

Mukul Pandya

The Wharton School,
University of Pennsylvania

Artificial intelligence, or AI, is transforming every aspect of business and society, including the very nature of customer engagement. This point of view is shared by several authors of books about AI and marketing who believe that AI will dramatically change customer engagement in the future, at least in the following seven ways:

- First, through the use of AI, companies will be able to redefine customer experience by making it possible to personalize products and services at scale.
- Second, AI will help companies to move away from short-term transactional interactions with their customers and build long-term relationships that can help

them maximize customer lifetime value (CLV).

- Third, AI will allow firms to offer more immersive, impactful experiences to their customers, especially in conjunction with emerging technologies such as the metaverse.
- Fourth, AI will sharply lower the cost of making predictions and enable companies to use new techniques to build customer engagement.
- Fifth, AI will make companies more mindful that, if they do not consciously introduce new experiences, customizing user interactions can trap people in silos.
- Sixth, AI will help companies make better and smarter use of data to improve customer engagement.
- Seventh, AI will transform customer engagement by changing what products and services are offered to customers and also how companies create and deliver these products and services.

Understanding the impact of AI on customer engagement is crucial for the entire C-suite. It matters, of course, to marketing managers who are at the front lines of the battle to win customers, but it is equally important for leaders who are involved with business strategy, R&D and innovation, and finance. Understanding customer engagement can also shed light on human relations, which is central to human resources and talent management. These books are valuable guides for executives in each of these roles.

1. **Personalize products and services at scale:** As the chief measurement strategist at Google, Neil Hoyne has led thousands of projects for companies that want to engage effectively with customers.[1]

Asked to suggest a book that will help readers understand how AI is reshaping customer engagement, he recommended *The AI Marketing Canvas: A Five-Stage Road Map to Implementing Artificial Intelligence in Marketing*[2] by Raj Venkatesan,[3] who teaches marketing and analytics at the University of Virginia's Darden School of Business, and Jim Lecinski,[4] a clinical associate professor of marketing at the Kellogg School of Management at Northwestern University. "Jim spent twelve years at Google leading these types of efforts for customers. You're not going to find many authors with that level of depth intersecting both research and practice," Hoyne notes.

In *The AI Marketing Canvas,* Venkatesan explains that AI can "supercharge the customer experience through personalization at scale." Explaining this philosophy in *Forbes* in August of 2021, Venkatesan spoke about Peter Chang, a Chinese restaurant that he often visits near the Darden campus. The restaurant staff greet him by name, know where he likes to sit and what he likes to order, and are familiar with his family.

In other words, for the staff at Peter Chang, Venkatesan is not a nameless, faceless customer; he is a human being with whom they have built a relationship. That allows them to offer him a personalized service. "Brands like Chick-fil-A and Starbucks are trying to achieve this personalization using AI in their mobile apps and translating the insights from the app into personalized in-store customer experiences also," Venkatesan said to *Forbes*.

AI, with its ability to crunch vast amounts of data at phenomenal speed, will allow companies to apply this degree of granular personalization to every aspect of customer engagement, according to Venkatesan. This includes acquisition, retention, growth, and advocacy, which is the term he uses for word of mouth. Venkatesan and Lecinski recommend a five-step progression for top executives who want to use AI to build customer engagement.

- Foundation: lay the foundation by collecting the first-person data that can be used to train the AI algorithm.
- Experimentation: conduct experiments to learn how to personalize a single aspect of customer engagement.
- Expansion: personalize more than one aspect of the company's engagement with the customer.
- Transformation: every aspect of customer engagement is personalized.
- Monetization: use AI to generate a fresh revenue stream.

As Venkatesan points out, companies in industries ranging from consumer packaged goods to banking and retail now use AI to personalize engagement with customers.

2. **Build long-term relationships to maximize customer lifetime value (CLV):** Hoyne's *Converted: The Data-Driven Way to Win Customers' Hearts*[5] has several nuggets of information about the use of AI to build customer engagement. Though he has written this book in his personal capacity, Hoyne draws extensively upon his experience at Google and as a senior

fellow with Wharton Customer Analytics at The Wharton School of the University of Pennsylvania.

Asked how brands can become better at using AI for mass-personalization at scale, Hoyne says, "When we talk about personalization, we have to broaden our lens a little bit to say, what do we know about these customers and what are their actual needs? How do we deliver them? That also removes that pressure from us where we don't need to get them to convert every time."

Hoyne offers two examples from the travel industry to make his point. He often books hotels using an online travel aggregator or an OTA, which generally buys its inventory of hotel rooms from other sites. "If they're focused only on short-term transactions, their goal is, every time I go to the site, to get me to book a new hotel room. That's where their investments go, their marketing messages, and their call to action. So they'll underinvest in personalization to improve the customer experience, because in their mind it doesn't immediately lead to any more incremental reservations," he notes. "Instead, if they took a step back to look at the lifetime value of customers and provided them with ancillary benefits, they would see that those customers stick around longer and spend more money."

Another example is that of a mobile app developer that was managing an airline app. The company tracked the number of people who used the mobile app to immediately book a new reservation. "That's what they were focused on, ROI (return on investment)," Hoyne notes.

The company discovered that the lifetime value of people who used the app compared to that of those who did not was roughly the same, so the mobile app was not adding anything incremental to the experience.

They started adding more features to help customers with their reward points, navigate their way around the airport, and manage their travel reservations. This did not lead to immediate bookings, but they found that customers were happier and made more bookings, thus increasing their lifetime value.

3. **Create more immersive experiences with AI and the metaverse:** Rajesh Jain, founder of Netcore, one of India's most innovative digital marketing companies, recommends *The Metaverse: And How It Will Revolutionize Everything* by Matthew Ball.[6] Jain notes that "The book gives a glimpse about how brands and customers will engage in the future – in a world of infinite computing power, storage, bandwidth, consumer data, and AI. The future is about how brands can build better hotlines with their customers, and the metaverse will be one of the key places to engage with present and future customers. Every brand will need a metaverse strategy and Matthew Ball's book serves as an excellent introduction to get started."

Ball is the managing partner of the early stage investment fund Epyllion and the former global head of strategy for Amazon Studios.[7] He adapted his book into an essay that appeared as a cover story in *Time* magazine in July 2022.[8] In January 2022, Ball

notes, "Microsoft paid $75 billion for the gaming giant Activision Blizzard, which will provide building blocks for the metaverse." Ball quotes an estimate by McKinsey and Company that "corporations, private equity companies, and venture capitalists made $120 billion in metaverse-related investments during the first five months of this year."

What does this have to do with customer engagement? As a case in point, consider Roblox, an online game platform and game creation system that lets users develop games and play games that other users have created. In January 2022, Roblox was averaging more than 4 billion hours of use a month. "Part of Roblox's surging engagement is driven by its growing user base," Ball writes. "From Q4 (fourth quarter) 2018 to January 2022, average monthly players increased from an estimated 76 million to more than 226 million (or 200%), while average daily players grew from around 13.7 to 54.7 million (or 300%)."

Though it might seem that the games that Roblox users play have little to do with AI or customer engagement, the metaverse is likely to usher in the next generation of the Internet. Companies will be able to use similar technologies such as virtual or augmented reality to allow customers to try on virtual garments before ordering them, visit a digital replica of a home before making an offer, or test drive a virtual car before buying the real one, and much more. Some companies such as Nike, Samsung, and Hyundai have already begun using AI

in conjunction with these technologies to deepen customer engagement.[9]

4. **Make precise, cost-efficient predictions:** In many aspects of business, AI makes excellent predictions cost-efficiently. This is the principal argument of the book *Prediction Machines: The Simple Economics of Artificial Intelligence,*[10] recommended by Kartik Hosanagar, a professor of operations, information, and decisions at The Wharton School of the University of Pennsylvania and faculty co-lead of the AI for Business initiative.[11]

Authored by Ajay Agrawal,[12] Joshua Gans,[13] and Avi Goldfarb,[14] economists at the University of Toronto's Rotman School of Management, this book posits that the rise of AI has reduced the cost of making predictions, just as the coming of the World Wide Web in the mid-1990s reduced the costs of search, communication, and so on.

"The central point of the book is that from a business/economic standpoint, AI increases the accuracy and reduces the cost of making data-driven predictions. The rest of it is focused on managerial implications of that," Hosanagar says. "For example, if predictions are cheaper and accurate, the premium might then go towards how one uses those predictions for judgment and managerial decisions. While this isn't specifically on customer engagement, arguably all managers should care about this."

According to Jerry Wind,[15] professor emeritus of marketing at The Wharton School, and editor of this special issue, AI can be particularly valuable when companies want "to focus on the need to predict levels of customer engagement and the determinants of engagement."

At a talk about the book at Google in 2018, Goldfarb pointed out that advances in AI have largely been the result of advances in prediction technology, which enables better, faster, and cheaper predictions.[16] Why does that matter? When computing reduced the cost of arithmetic, new applications for computing proliferated. A similar process is at work with the falling cost of prediction driving an upsurge in AI. Prediction means using information you have to fill in information you do not have. The first predictive AI apps addressed loan defaults, helping banks to figure out whether their loans would be repaid.

The insurance industry also loves AI tools because it is steeped in making predictions. Insurers want to know if a customer is likely to make a claim or not, and for how much. New ways of thinking about prediction have arisen because the cost of prediction has fallen.

For example, medical diagnosis is, at heart, a prediction problem in which a doctor looks at apparent symptoms and fills in the missing information about the cause. Autonomous driving also solves a prediction problem. Automakers try to predict what a good human driver would do under the prevailing circumstances on the road. Once they can do that, they can build AI-powered cars to do what a human driver would. The key element in understanding and identifying new opportunities for cheap predictions is filling in missing information.

According to Goldfarb, when coffee becomes cheap, consumers will drink more of it. That is Economics 101. As customers consume more coffee, though, the demand for complementary products such as cream and sugar also increases. So next, companies should strive to find those complementary products.

As the cost of prediction falls, what will become more valuable? By answering that question, we should be able to discover where AI could make the most difference. For example, one complementary product is data. Data is often described as 'the new oil' because the falling cost of making predictions through AI has vastly increased the demand for data. Among myriad other things, data is used to fuel insights that will help all top executives to understand the impact of AI on customer engagement.

5. **Understand the limits of customer choice:** In his own book, *A Human's Guide to Machine Intelligence: How Algorithms Are Shaping Our Lives and How We Can Stay in Control,* Hosanagar devotes a chapter to AI and customer engagement.[17] "My book covers how AI is driving consumer choice," he says. "But it also warns that when personalization systems overly rely on past data and don't explore enough, it can lock customers in filter bubbles where they see more and more of what they have already liked and consumed at the expense of discovering new items and media and expanding their tastes." According to Wind, this represents a tremendous "opportunity for those using

AI to design the consumer engagement, to move beyond historical data and augment it with new experiences."

Hosanagar explains, "all of us realize how much of our lives are shaped by decisions we make online." This may be through searching for products or services on Google, connecting with friends on Facebook, or shopping for books or other products on Amazon. "Many of us are aware that the companies running these sites are guiding our choices by customizing our experiences. Personalization algorithms help us choose the optimal products to buy on Amazon, the best movies to watch on Netflix, the ideal person to date on Tinder or Match.com, or the most useful contacts on LinkedIn."

Hosanagar notes that though we might imagine that we are making these choices, the fact is that the websites exclude a lot of the choices they offer to customers and present a relatively narrow set of options from which to make choices. As a result, some 80 percent of viewing hours streamed on Netflix originate from automated recommendations presented by the site's AI-powered personalization algorithms. "Products are often designed in ways that make us act impulsively and against our better judgment," he writes. "Popular design approaches such as the use of notifications and gamification to user engagement exploit and amplify human vulnerabilities."

6. **Reinvent customer engagement through data:** If data has indeed become the new oil it may be wise to learn how to drill

for it. That is what the book, *Data Driven: Harnessing Data and AI to Reinvent Customer Engagement*, claims to do.[18] The authors are Tom Chavez,[19] the CEO of super{set}, who built a data management platform called Krux that Salesforce acquired in 2016; Chris O'Hara, a former columnist; and Vivek Vaidya, a serial tech entrepreneur and former co-founder and CTO of Krux. This book recommendation comes from Greg Shea,[20] an adjunct professor of management at The Wharton School.

According to the authors, "Ubiquitous connectivity only partially explains the rise of Internet giants and the surprising success of digital upstarts, such as Spotify, Tinder, and Twitch. A powerful steel thread runs through all these companies. What they have in common is data: the ability to capture it – increasingly every scrap – and put it to work to generate insights, recommendations, and offers that dazzle their customers."

The authors write that anyone who is in the business of engaging with existing customers or trying to find new ones can "run but you can't hide from data. It is the fuel that enables any company to know its customers intimately, improve its products, deliver better customer service, optimize any business process, and predict the future direction of markets."

7. **Transform the way products and services are delivered:** Saikat Chaudhuri,[21] faculty director of the Management, Entrepreneurship & Technology Program and the Berkeley Haas Entrepreneurship Hub at

the Haas School of Business of the University of California, Berkeley, suggests *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World,*[22] by Marco Iansiti[23] and Karim Lakhani,[24] who are on the faculty of Harvard Business School.

"I recommend this book because many manuscripts on AI these days focus heavily on the technology and possibilities that AI offers," says Chaudhuri. "However, the deployment of AI can fundamentally alter not just what products and services are offered to customers, but how organizations should create and deliver them. This requires product, process, and business model innovation. Iansiti and Lakhani tackle how firms can reinvent themselves to take advantage of the opportunities on the horizon and navigate the challenges as well as the risks, suggesting new architectures that can help us enhance scale, scope, and learn. They bring together technology, strategy, and organization in provocative ways that offer profound food for thought."

In this book, business leaders will find several instances of how AI-driven companies are challenging traditional industry leaders in customer engagement. Consider retail. The authors show how Amazon, with its AI-powered processes, is taking on Walmart. They write: "No industry is feeling the impact of Amazon more keenly than retail. Amazon's convenience, low prices, personalization, and recommendation capabilities and software-enabled logistics infrastructure was a formidable challenge to traditional firms. In 2017, we saw more than twenty

long-standing retailers file for bankruptcy, and in 2018, even the 105-year-old giant Sears joined this list. Walmart – the world's largest company by revenue – is doing all it can to avoid that fate."

The authors write that Walmart has hardly been shy about investing in technology. "For decades, it set the standard in retail supply chain technology and network infrastructure, with its constantly evolving retail link system and its early commitment to EDI and RFID technologies. A data rich supply chain has consistently been an important part of Walmart's operating model, and the key to its impressive scale," they say. Still, for all its technological prowess, Walmart has had to transform itself to deal with Amazon's challenge. "To put up a credible fight with Amazon, Walmart is re-architecting its operating model on a digital and AI-enabled foundation," they say.

For leaders who want to understand how AI will shape the future of customer engagement, there does not appear to be a single blueprint showing them what kind of organization to build or what kind of AI system to develop to best engage customers and other stakeholders. The wide range of perspectives that these books and their authors provide will nonetheless help to prepare them for whatever the future may hold. ■

Author Bio

Mukul Pandya is the founding editor in chief and executive director of Knowledge@Wharton, the online research and business journal of the Wharton School. After retiring from K@W, Mr. Pandya was a senior fellow with Wharton Customer Analytics and AI for Business. A four-time award winner for investigative journalism, Mr. Pandya has published articles in *The New York Times, The Wall Street Journal, The Economist, Time, The Philadelphia Inquirer,* and more. He has written or coauthored four books.

Endnotes

1. https://executiveeducation.wharton.upenn.edu/faculty/neil-hoyne/
2. https://www.sup.org/books/title/?id=32597
3. https://www.darden.virginia.edu/faculty-research/directory/rajkumar-venkatesan
4. https://www.kellogg.northwestern.edu/faculty/directory/lecinski_jim.aspx
5. https://www.penguinrandomhouse.com/books/688629/converted-by-neil-hoyne/
6. https://wwnorton.com/books/9781324092032
7. https://www.matthewball.vc/about
8. https://time.com/6197849/metaverse-future-matthew-ball/
9. https://influencermarketinghub.com/metaverse-brands/
10. https://store.hbr.org/product/prediction-machines-updated-and-expanded-the-simple-economics-of-artificial-intelligence/10598
11. https://oid.wharton.upenn.edu/profile/kartikh/
12. https://www.rotman.utoronto.ca/FacultyAndResearch/Faculty/FacultyBios/Agrawal
13. https://www.joshuagans.com/
14. https://www.rotman.utoronto.ca/FacultyAndResearch/Faculty/FacultyBios/Goldfarb.aspx
15. https://marketing.wharton.upenn.edu/profile/windj/
16. https://www.youtube.com/watch?v=ByvPp5xGL1I
17. https://wsp.wharton.upenn.edu/book/a-humans-guide-to-machine-intelligence/
18. https://www.mhprofessional.com/data-driven-harnessing-data-and-ai-to-reinvent-customer-engagement-9781260441536-usa
19. https://tomchavez.ceo/
20. https://www.gregoryshea.com/
21. https://haas.berkeley.edu/faculty/saikat-chaudhuri/
22. https://hbsp.harvard.edu/product/R2001C-PDF-ENG
23. https://www.hbs.edu/faculty/Pages/profile.aspx?facId=6482
24. https://www.hbs.edu/faculty/Pages/profile.aspx?facId=240491

©shutterstock.com/BRO.vector

Mars' ACE: Using AI and Behavioral Data in Ad Testing with High Correlating Sales

Laurent Larguinat describes Mars' ACE, a proprietary AI tool which analyzes large sets of consumer behavioral data to predict the sales effectiveness of the company's adverts far more effectively than traditional methods.

Laurent Larguinat
Mars, Inc.

Advertising testing is a traditional method used by marketers to gain insights into a consumer's response to ads. Researchers typically use standard survey methods in which respondents are asked about ad recall, ad preference, or desire to purchase in future before the ad is launched.

But traditional ad testing is limited by its declarative survey approach. We need an alternate method to drive sales and create a sustainable competitive advantage. Mars developed its own proprietary method of assessing effectiveness by using consumer

The result is an 85 percent accuracy in predicting whether ads are likely to lead to a sale.

behavioral methods and applying artificial intelligence (AI).

We used AI algorithms to analyze behavioral data, including eye tracking, facial reactions, skin conductivity, EEG, and more, and ultimately link benchmarks to sales success. The result is an 85 percent accuracy in predicting whether ads are likely to lead to a sale.

The Mars approach

Mars has been proudly family-owned for over a century. Having the independence of a privately held company allows us to think along a timeline spanning generations, not fiscal quarters. Guided by enduring principles, we can invest in the long-term future of our business, our people, and the planet. Our longevity gives us the freedom to try new things, developing our own tools to forge new paths.

In the consumer packaged goods (CPG) industry, Mars is widely recognized for its strong insights culture, an evidence-based approach to marketing decisions, and a reluctance to used biased or unscientific approaches to understanding consumers. For more than fifteen years, Mars focused its media research on the best behavioral tools. We used industry-leading single source panels and natural experiments to measure the direct sales impact of media exposures at the household level.

In a world of data-driven marketing, measuring sales impact is important but not fast enough to drive agile business decisions. To achieve extreme agility, we developed a behavioral proxy measure of business performance. We call it Agile Creative Expertise, or ACE.

ACE and sales prediction

ACE is an ad-testing tool powered by AI that emphasizes the emotional response of a test audience to a video ad while also weighing traditional metrics such as skip and click rates. The viewers, who were recruited by Mars partners and are compensated, consented to being tracked. We anonymized and aggregated their data before sharing the results with Mars. In designing ACE, we tested a few hundred pieces of content for which we already had a clean and reliable measure of market sales performance (single source). For each of these pieces of content we tracked attention, emotion, views and skips, and sales, all second-to-second.

To this information we added the data we collected through the largest ever CPG behavioral science study, which we conducted with Nielsen. In it we tested 300 creative executions, or marketing components, on participants in seven countries. We applied AI on this combined dataset to understand if and how the emotion, attention, and views predicted sales impact. The result was an algorithm that is nearly 85 percent predictive.

In 2019, AI models were evolving, enabling emotional recognition directly from a mobile device. This development helped the technology to jump from the lab into the living rooms, kitchens, and cars of consumers. After more than six years of testing and retesting AI-fueled advertising research approaches, Mars Horizon, our insights team, launched ACE in 2020.

Validation: considering other solutions

To ensure ACE's credibility, we had to link it to predicting sales. Without this confirmation, the company would have rejected any new approach to advertising research.

To this end, we rated our content between one star (no impact on sales) and four stars (more than 15 percent impact on sales and in the top quintile of all content in a given category). Before ACE, about 40 percent of all the content we developed was rated three or four stars. ACE increased this number to more than 70 percent.

To ensure ACE's credibility, we had to link it to predicting sales.

We were able to prove that ACE correctly predicted consumers' sales response to more than 4,000 ads, making us confident that our facial response data offers a proxy metric for sales and, ultimately, the effectiveness of the ad. Our monitoring of the correlation between facial response data and sales is, of course, ongoing, and allows us to correct our course and benchmarks as our volume of data increases.

This approach to testing is proving to be better at predicting in-market success than the more traditional methods we once used. It is also faster and more affordable.

We also investigated a number of more invasive technological solutions, including skin conductance and EKG, but found that these methods are more complex and do not provide any additional information about consumers' attention and engagement.

Deployment and adoption

The Mars company has now widely embraced ACE. In 20 months, we completed more than 1,000 studies in over 25 countries. Our results helped Mars to improve its digital adverts on TikTok, Facebook, Instagram, YouTube, Amazon, and more, generating an additional $30 million in sales.

Today, we use ACE to test ads for all our brand's digital channels. We proved that an ad with a higher attention score will travel across media channels better, driving sales. We also use attention to understand

which components of ads are winning consumers' minds, and which are just confusing them.

The future of AI in ad testing

Through this AI approach, we ensure that our ads inspire attention and emotion, guaranteeing that we increase consumer engagement.

Ultimately, advertising is a mix of science and art. And art is difficult to evaluate without human input. Consumers will always play an important role in future advertising research. At Mars, we are excited to continue producing adverts and digital content that resonate with consumers.

We believe that AI will only become more valuable to ad testing. As tools that replicate real consumer behaviors proliferate, using AI for eye tracking, determining attention, and more, marketers will be better able to optimize ads. Perhaps one day, the AI itself will optimize the ads on the fly. ◼

Author Bio

Laurent Larguinat leads the Mars Horizon team, responsible for marketing and sales sciences at Mars Inc. The team collaborates with world-class academics, research organizations, and start-ups to solve the company's marketing and sales challenges.

Machines that Dream:

How AI-Human Collaborations in Art Deepen Audience Engagement

Refik Anadol
Artist

Pelin Kivrak
Refik Anadol Studio
and Tufts University

Human engagement in the realm of digital art will take many forms, both in the physical world and in the metaverse. Refik Anadol and Pelin Kivrak explore how Refik Anadol Studio experiments and collaborates with AI to create art that evokes multiple senses, using immersive experiences to engage with audiences.

The metaverse, with its immersive virtual reality environment, promises a dynamic canvas for the bold and imaginative. It is a new territory in which to explore our desire for collaborations between machine and human, whether by building a new virtual universe with its own intrinsic values or a parallel world where presence and perception have deeper meanings.

Using data as pigment

Refik Anadol Studio (RAS) in Los Angeles has been experimenting with hybrid forms of AI-based art-making, combining the creative possibilities of the metaverse with the studio's decade-long vision of embedding new media arts in architecture. Anadol's renowned body of works takes publicly available datasets which represent the human experience in various ways, ranging from millions of photographs of bustling urban centers to a comprehensive image and video dataset capturing the glaciers of the world.

The primary thread that runs throughout the studio's visualizations of the unseen world is the use of data as pigment to create enriched immersive environments. Taking the data that surrounds us as the primary material and the neural network of a computer as a collaborator, the studio's site-specific artificial intelligence

(AI) data paintings and sculptures, live audiovisual performances, and environmental installations encourage us to rethink our engagement with the physical world, decentralized networks, collective synesthesia, and the quantifiable healing power of art.

If machines can learn or process individual and collective memories, can they also dream or hallucinate about them?

Since 2012, RAS has been conducting interdisciplinary research on the interconnection of the human mind, architecture, and aesthetics in an effort to answer this question: If machines can learn or process individual and collective memories, can they also dream or hallucinate about them? After experimenting with digital paintings and sculptures of architectural data and exhibiting data-based immersive projections in public spaces, RAS started working with AI programs and machine learning algorithms using a diversity of data, be it visual, auditory, seismic, geographic, cultural, or institutional.

An ambitious beginning

The studio's journey toward this unique machine intelligence approach to creating art began with the 2016 project titled *Archive Dreaming*. Refik Anadol was invited to join the Google artist in residence program for Artists and Machine Intelligence (AMI). There, engineers and thinkers in the field, including Google's own Blaise Aguera y Arcas, Kenric McDowell and Mike Tyka, helped the entire RAS team to learn and use AI to create *Archive Dreaming*.

To complete the large-scale work, RAS collaborated with SALT Research collections in Istanbul, which provided access to 1.7 million documents from the 17th to 20th century archives of the Ottoman Bank, a dominant financial institution during the Ottoman Empire. We employed machine learning algorithms to search and sort relationships among these data points while translating interactions of the multidimensional data in the archives into an immersive and interactive media installation.

Archive Dreaming, presented in Istanbul in 2017 as part of *The Uses of Art: Final Exhibition*, supported by the Culture Programme of the European Union, was user-driven. Yet when idle, the installation 'dreamt' of unexpected correlations between documents, yielding generative, aesthetic, and fresh visuals
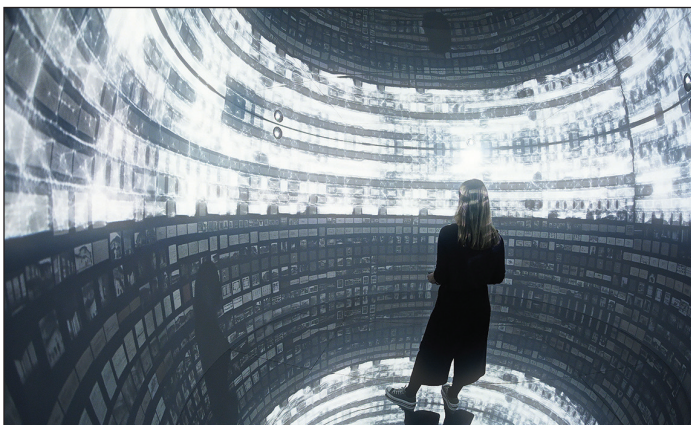
Archive Dreaming (2017), Istanbul, Turkey

representing those serendipitous links and overlaps.

Refik Anadol Studios translated the resulting high-dimensional data and interactions into an architectural immersive space, projected onto a circular area in 7680 × 1200 resolution. The installation – at times displaying swirling points of light in a rotating tunnel, or flashes of brightness with shapes that shimmer and disappear, or cloud clusters and showers of illumination – allowed the audience to browse the millions of documents and images within the archive. Curious art lovers could peak inside this universe of data and zoom in on any file at their pleasure. After its debut at SALT Galata, an extension of the project was exhibited at the Ars Electronica Festival 2017 in the section titled *Artificial Intelligence.*

RAS rendered these images as three-dimensional forms determined by the similarities between documents found by machine learning algorithms. Thousands of Istanbulites and researchers from around the world visited the installation in the heart of bustling Istanbul.

Kenric McDowell, who leads Google's AMI program, described the generative potential that this groundbreaking project offered at the intersection of culture and technology: "The experience is visually and architecturally impressive, but what's most exciting about it is the way it reframes the archive as a multidimensional space of images connected by features, that is to say, visual relatedness. The piece also uses the machine learning model to hallucinate new images that *could* exist in the archive but don't. To my knowledge this is the first time a museum archive has become generative in this way."[1]

Architectural 'consciousness'
Another turning point in the studio's visualizations of institutional ar-
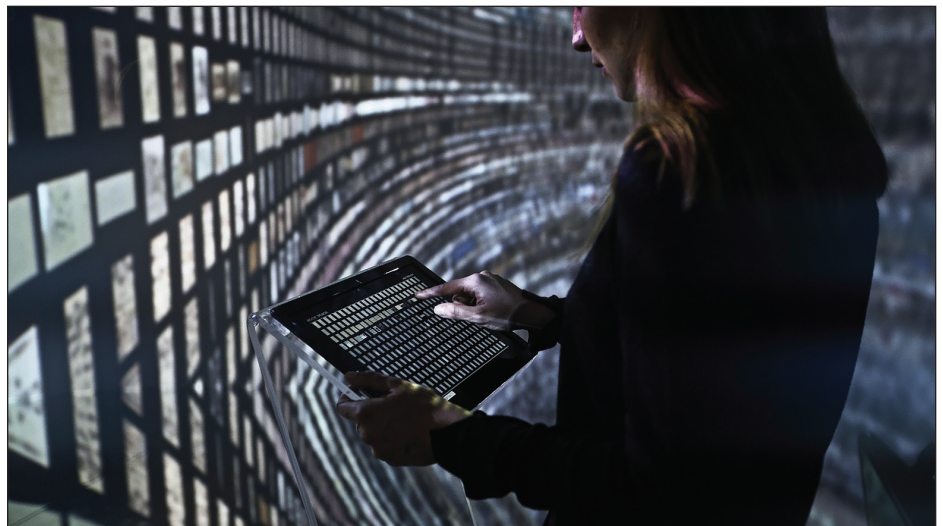
In addition to providing an aesthetic experience, *Archive Dreaming* offered user interaction with the digitized archive.

Walt Disney Concert Hall Dreams (2018), Los Angeles, USA

chives was *Walt Disney Concert Hall Dreams (WDCH Dreams)*, a collaboration with the Los Angeles Philharmonic (LA Phil) to celebrate its 2018/19 centennial season.

For this project, RAS developed a unique software to process the orchestra's digital archives – 45 terabytes of data containing 40,000 hours of audio from 16,471 performances. The AI parsed the files into millions of data points that it then categorized by hundreds of attributes using deep neural networks with the capacity to both record the totality of the LA Phil's memo-

ries and create new connections between them.

This 'data universe' generated something new in image and sound by awakening the metaphorical consciousness of Walt Disney Concert Hall. The visual art was projected onto the building's curved exterior walls, a swirl of moving images ranging from a rainfall of data, to musicians playing, to waves of white blocks surging like a digital ocean, and other abstract imagery.

The result was a radical visualization of the LA Phil's first century and an exploration of the synergies
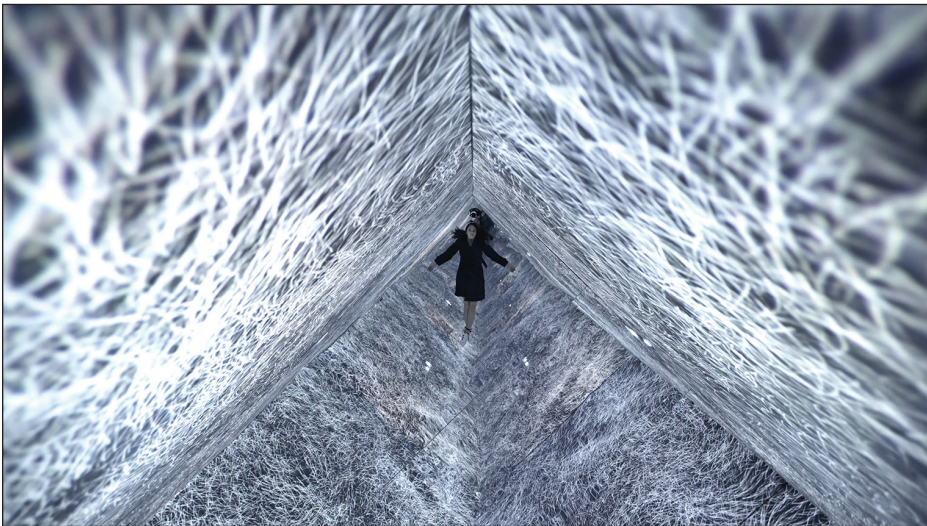
Refik Anadol Studio's *Infinity Room* has been physically experienced by more than 2 million people around the world

Quantum Memories, (2021), National Gallery of Victoria, Melbourne

Machine Hallucinations: Nature Dreams, (2021), Koenig Gallery, Berlin

between art and technology, architecture, and institutional memory. The work presented the philharmonic's entire digital archives in a non-linear way. It also featured an interactive companion installation in a U-shaped room in which two-channel projection provided multiple experiences for visitors.

RAS also collaborated with the Philadelphia Orchestra in the Spring of 2022 on a project titled *Beethoven: Missa Solemnis 2.0*. In Verizon Hall, a new AI-based artwork unfolded for an audience as they listened to the composer's 1823 masterwork. The research question for this unique artwork's dataset was: "Why not use artificial intelligence [AI] to try to reconstruct the reality of what Beethoven could imagine?" With this motivation, the RAS team started compiling a dataset of 12 million images of buildings that Beethoven could have encountered in Europe during his life. We then used this archive to train a custom algorithm that generated a new data universe of artificial architectural images for us to curate into the final artwork.

The dynamic visual artwork, consisting of the machine's hallucinations of alternative European architectures, alluded to religious spaces and iconography, generating

the feel of a virtual cathedral for the performance of this sacred music piece. We created another software tool that allowed the AI to listen to the orchestra performing Beethoven, simultaneously generating and projecting its dreams for the audience. This addition further manifested the vision of the borderlessness of art that the studio espoused with the *WDCH Dreams* project.

These early works of interdisciplinary and international collaboration paved the way for the global impact of RAS' immersive artworks to gradually expand. *Infinity Room*, an artwork developed in 2015, combines the boundlessness of space with the endless permutations of machine intelligence, inviting visitors to step into a 360 degree mirrored room that uses light, sound, scent, and technology to offer a seismic perception shift.

Originally presented in collaboration with the 2015 Istanbul Biennial, *Infinity Room* began traveling the world. More than two million people have now experienced it. In 2021, the National Gallery of Victoria exhibited *Quantum Memories*, the studio's epic scale investigation of the intersection of architecture, machine learning, and the aesthetics of probability on the largest LED screen it had ever used. The exhibit has welcomed 1.7 million visitors, the largest audience to experience a digital artwork in Australia.

Machine Hallucinations: Nature Dreams, a solo exhibit of AI data sculptures and paintings based on nature-themed datasets in Berlin's Koenig Gallery, drew 200,000 visitors in five weeks, the largest audience to visit a gallery in Europe. *Machine Hallucinations: Space* was the product of our collaboration with with NASA's Jet Propulsion Laboratory, beginning in 2018. It is an immersive art piece rooted in publicly available photographs of space taken by satellites and spacecraft. Ninety thousand visitors in

Details from *Living Architecture: Casa Batlló*

The live projection mapping performance of *Living Architecture* on the façade of Casa Batlló on May 7, 2022 has been viewed by an enormous crowd of 50,000.

three weeks gave it the largest audience ever to view an artwork in Hong Kong.

Engaging a diverse audience

RAS's broad experience of creating such large-scale, immersive, multi-sensory installations around the world motivated the team to take these ideas to the metaverse. In the spring of 2021, RAS had the opportunity to explore its potential for creative production with a collaborative artwork centered around Barcelona's iconic 1906 Gaudí building, Casa Batlló.

We created *Living Architecture* as an example of how an artwork could engage a diverse audience by existing in multiple and dynamic virtual forms while still exhibiting a strong connection to the physical world. This long-term project began when RAS was commissioned to create *In the Mind of Gaudí*, an AI-based immersive experience presented in an LED cube room, lined with six screens, in the basement Casa Batlló.

For this initial 360 degree experience, RAS collected approximately one billion images consisting of Gaudí's sketches, visual archives of the building's history, academic archives, and publicly available photos of Casa Batlló found on various Internet and social media platforms. We processed them with machine learning classification models which sorted images into thematic categories.

With the help of our custom generated software and fluid simulation models, making the images dynamic, we then transformed this data universe, inspired by Gaudí, generated by machine, and curated by humans, into an AI data sculpture. We used subsets of the data to train an AI model, causing the machine to hallucinate new aesthetic images and color combinations through algorithmic connections. We then clustered the images into thematic categories to better understand the semantic context of the data universe.

This expanding data universe was not just the interpolation of data as synthesis, but a latent cosmos in which the hallucinations of the AI was the main creative and artistic currency. To capture these hallucinations, we used NVIDIA's Style-GAN2 adaptive discriminator augmentation to generate an AI model through which the machine could process the archive. The model was trained on subsets of the sorted images, creating embeddings in 1024 dimensions.[2]

> RAS collected approximately one billion images consisting of Gaudí's sketches, visual archives of the building's history, academic archives, and publicly available photos of Casa Batlló.

The next stage is the data pigmentation. We combined visual elements from a range of sources into single images. For more than ten years, we have been experimenting with custom software and fluid simulation models to give these visuals a unique dynamism. Fluid simulation allows a computer to emulate and generate the visual qualities and behavior of a fluid. It is one of our studio's signature visual effects in data visualizations. We have used it at various levels of complexity including real-time and interactive animation.

Redefining the museum experience for visitors

For *Living Architecture: Casa Batlló*, we synthesized the vast data from Gaudí archives into ethereal data pigments and eventually into fluid inspired movements. *In the Mind of Gaudí* was part of Casa Batlló's monumental re-opening after the pandemic, a series of immersive and interactive tours that created new connections between the masterpiece and its visitors.

The installation used a 360 degree art experience in the LED cube, billed as the world's first 10D experience, to offer unprecedented insight into the Catalan architect's mind, using the world's largest digital Gaudí library.[3] The technology that we used to represent Gaudí's eccentric aesthetic, a dynamic multi-sensory space with more than 1,000 digital projections, twenty-one audio channels, binaural sound, and scent, redefining the museum experience and attracting audiences from around the globe and from all ages and backgrounds.

We went on to build a second artwork from the same Gaudí archive, one that can be experienced and celebrated both physically and in the metaverse as an NFT (non fungible token). NFT artwork is digital art that can't be duplicated, and therefore has value as a unique item. In *Living Architecture*, RAS radically interpreted Gaudís building and its memories, making it the first UNESCO World Heritage Site to be represented in a dynamic NFT.

Using the building's iconic façade as a canvas, we displayed an algorithmically generated, dynamic abstraction. The digital piece was also shown at Rockefeller Plaza in Manhattan, New York. To amplify the visual and physical impact of the work, we created a custom scent and audio soundtrack for the pre-sale of the exhibition which complemented the urban environment.

The work was sold at Christie's 21st Century Evening Sale on May 10, 2022 for $1.38 million. We donated 10 percent of the proceeds to institutions that work with neurodiverse adults and children. Three days prior, we had projected a mapping version of the piece on the facade of Casa Batlló before 47,000

people. *Living Architecture* embodied how advanced our experimentation with machine learning tools and blockchain technology have become; we used sensors in Barcelona that collect real-time environmental data which caused the NFT piece to change along with the city's weather and the art projected on the façade.

An ongoing experiential artwork

Technologist and startup CEO Sagi Eliyahu defines customer engagement as an ongoing emotional relationship that can best be described as the sum of multiple moments or the customer's overall emotional connection that results from the totality of the experience. The Casa Batlló project offered exactly this connection, culminating in the final sale of the NFT artwork.[4]

As Gary Gautier, Casa Batlló's director, put it, "The landmark was seen around the globe thanks to an unprecedented audiovisual show that was watched from Barcelona, New York, the metaverse, and thousands of houses, live. Those who had the chance to watch it, we watched a unique moment in the art history: bringing together a World Heritage site and the digital artwork from Refik Anadol."[5]

The ultimate NFT project perfectly reflected RAS's vision of reaching a diverse audience

> We see such meaningful collaborations between machine and human not only as aesthetic representations but also as productive experiments.

through multiple channels of sensory engagement. We see such meaningful collaborations between machine and human not only as aesthetic representations but also as productive experiments that can be used to create multiple forms of psychological fulfillment and empowerment in both the psychical and virtual world.

The NFT sale allowed us to contribute to a worthy cause, raising further awareness of such work as the meaningful merging of science, technology, and the arts to address global issues. It also demonstrated the potential of metaverse architectures to reaching a broader audience through decentralized creativity, as well as the innovative use of blockchain currencies, to raise funds for charities, providing tangible solutions to diverse global problems. ■

Author Bios

Refik Anadol is an internationally renowned artist, director, and leader in the aesthetics of data and machine intelligence. He is a lecturer in UCLA's Department of Design Media Arts. Anadol's work places creativity at the intersection between humans and machines. It has been exhibited at the MoMA, Centre Pompidou-Metz, National Gallery of Victoria, Venice Architecture Biennale, Arken Museum, Hammer Museum, Dongdaemun Design Plaza, Ars Electronica Festival, Istanbul Design Biennial, and more.

Pelin Kivrak is a writer, a senior researcher at Refik Anadol Studio, and a scholar of comparative media studies. She is a lecturer in Tufts University's Department of English. Kivrak earned a PhD in comparative literature from Yale University and a BA in literature from Harvard University. Her academic work focuses on contemporary art and literature's engagement with the philosophical and political genealogies of the concept of collective responsibility.

Endnotes

1. Kenric McDowell and Ben Vickers, Kenric McDowell in conversation with Ben Vickers (Part 1), CURA.26, https://curamagazine.com/digital/kenric-mcdowell-in-conversation-with-ben-vickers-part-1/.

2. We achieve a 1000-dimension computer representation of an image by passing it through a neural network classifier known as ResNeXt.

3. The work is part of a multi-sensory experience that uses the visitor's five senses through volumetric projections, binaural sound, movement sensors and immersive spaces.

4. Sagi Eliyahu, "Council Post: How AI Can Change Customer Experience and Engagement," *Forbes* (Forbes Magazine, November 16, 2020), https://www.forbes.com/sites/forbestechcouncil/2020/11/17/how-ai-can-change-customer-experience-and-engagement/?sh=5a00bc37116f.

5. "Gaudí's Casa Batlló: first UNESCO World Heritage site to become live NFT," May 10, 2022, https://www.catalannews.com/life-style/item/gaudí-s-casa-batllo-first-unesco-world-heritage-site-to-become-live-nft.

Unlocking Deeper Insights into Customer Engagement Through AI-Powered Analysis of Social Media Data

P.K. Kannan
Robert H. Smith School of Business, University of Maryland

Yi Yang
Hong Kong University of Science and Technology

Kunpeng Zhang
Robert H. Smith School of Business, University of Maryland

P.K. Kannan, Yi Yang, and Kunpeng Zhang describe their method for using AI to map and analyze the structure of social media engagement, which spans thousands of brands in different categories. By using this method, managers can extract valuable information about customers, trends, ties to other firms, and impending opportunities or threats.

Brands use social media channels to engage with their customers by posting content informing them of new products and services, requesting feedback, and increasing sales through influencer marketing. Users interact and engage with these posts by liking, commenting on, or sharing the content.

Using artificial intelligence (AI), companies can analyze the social engagement structure that spans several thousands of brands in different product and service categories, anticipating trends and opportunities in the market and gathering intelligence on competitors. Managers can use this analysis to engage customers and grow their business.

In our paper, "A Deep Network Representation Learning of Social Engagement,"[1] we investigated companies' less common uses of social media: anticipating trends and opportunities in the market and gathering competitor intelligence. Specifically, we focused on how AI can help firms to understand the structure of social engagement and draw insights from it to engage potential customers.

A user could interact with any number of different brands on these sites and a brand could have similar interactions with any number of users. If we consider each interaction between a brand and a user to be a link between them, social media interactions can be characterized as a network of links between brands and users, highlighting the users who different brands have in common and the brands that different users have in common.

We define these nodes of brands and users, and their links, as comprising the social engagement structure on a social media platform. Firms can analyze this structure, using deep learning auto-encoders to process the network data and identify the competitive structures – those that exist within the brands' markets as well as those that cross the boundaries into other products and brands.

Relevance of the Social Engagement Structure

The information inherent in the social engagement structure on social media platforms can be useful to companies. At a minimum, it can help them to identify other brands with which they share common users and fans. These brands could be either substitutes or complementary.

An automobile company could identify its close competitors from the engagement structure, spotting other brands of automobiles that users also like or with which they engage. These brands share the

Beyond such obvious conclusions is a wealth of latent information that our AI unearths from the overall network and that can be immensely valuable to companies.

same group of users and so could be substituted for one another. A hotel that highlights an airline with which it shares many common users is pointing out a complementary link which could lead to opportunities for cross-promotion and brand tie-ins.

But beyond such obvious conclusions is a wealth of latent information which our AI unearths from the overall network and that can be immensely valuable to companies.

The social engagement structure allows us to identify product-market boundaries, that is, brands that compete within a market, and changes in those boundaries that occur over time. Identifying the strength of competition between brands within the product market can inform strategy for next-generation product design, product positioning, new customer acquisition, and pricing and promotion decisions.

As technology advances, however, the product-market boundaries themselves are changing. These changes open the door to competitive threats and business opportunities outside the previously narrow boundaries of markets. The social engagement structure will immediately reflect these changes, allowing companies to identify both threats and opportunities.

And product-market boundaries are extremely fluid. Not only do technological developments upend markets but new products also change their structure. Tesla, which initially entered the high-end automobile market with an innovative fuel technology, has since rolled out products for the lower-end market, changing competition in that market as well.

Firms also acquire other firms and enter new markets. Amazon crossed market boundaries when it acquired Whole Foods and entered the offline grocery product market. In many such situations, product-market boundaries are defined on the basis

of SIC (Standard Industrial Classification) and NAICS (North American Industry Classification System) industry classification codes which cannot provide adequate indicators for emerging threats and opportunities.[2] The social engagement structure, however, offers a more dynamic view as customers react to these changes, and thus provides more foresightful information on these fluid boundaries.

Our AI-based technique differs from prior methods in that it scrutinizes disaggregated individual relationships between users and brands within the social engagement structure. This distinction is important when the product market is not specified in advance because it captures relationships between a diversity of brands in a range of markets.

Identifying such latent relationships is the essence of our approach.

For example, if a user likes American Airlines and Marriott, while another user likes Delta and Marriott, the indirect relationship between American Airlines and Delta is captured through Marriott, which is in a different market. Identifying such latent relationships is the essence of our approach.

Data Collection

We collected the social engagement data from Facebook because it is one of the largest and most representative social network platforms. We began by using Socialbakers, a social media marketing website, to obtain a list of U.S. brands with the most followers. Public fan pages on Facebook are categorized into several groups on Socialbakers, including brands, celebrities, community, entertainment, media, place, society, and sport. For this analysis, we focused on brands. Every brand on Facebook

Figure 1. The Overall Approach
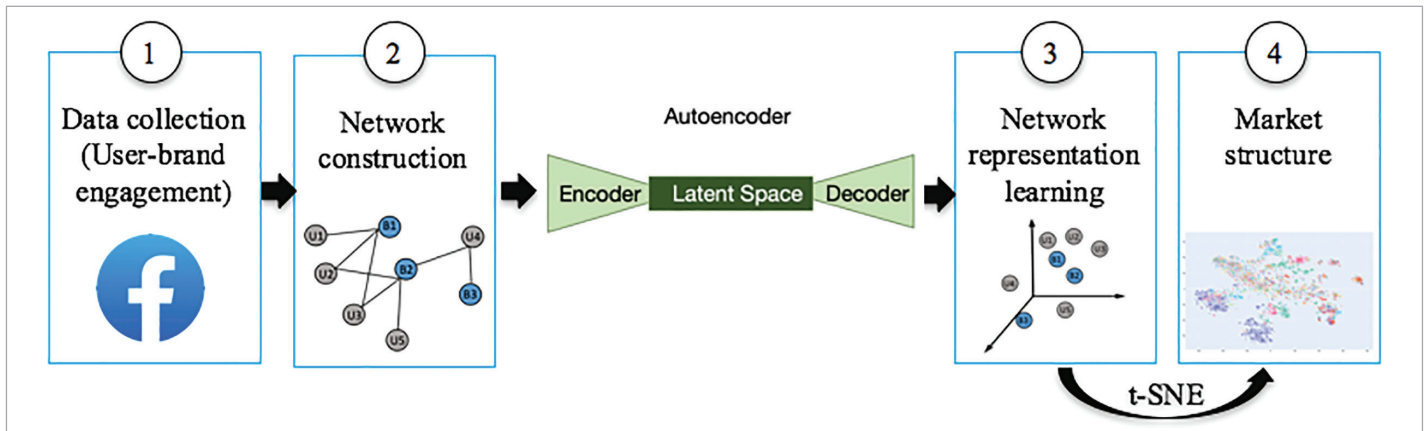
selects its category from a set of pre-defined options when it creates its public page. We included 5,478 different brands, spanning twenty-five categories.

We used Facebook Graph API to download from each brand page all visible activities including posts by the brand administrator as well as posts by users, including comments and likes on brand posts. To ensure privacy protection, we did not download any user profile information nor did we examine the content of user comments. All engagement activities were represented by unique user identifiers, regardless

of whether the user had a public or private Facebook profile, and by brand identifiers.

The data for this study covered the period between January 1, 2017 and January 1, 2018. In total, we used 106,580,172 user-brand engagement activities from 25,992,832 unique users.

To ensure data quality and robust results, we designed a set of rules to filter out fake users and their activities.[3] We then constructed a brand-user network including all of our selected brands and all users engaging with them. We considered a brand node and a user node to be connected if the

user engaged with the brand. The strength of a link between a brand node and a user node is defined by the frequency of engagement.

AI-Based Deep Network Representation Learning

The AI system we use is a network representation tool, also known as network embedding, whose goal is to compactly represent the underlying structure and relationships in a network.

The system uses an autoencoder whose mapping function translates the original network data of brands and users into a low dimensional form, preserving the brand

Figure 2. The Global Market Structure[7]
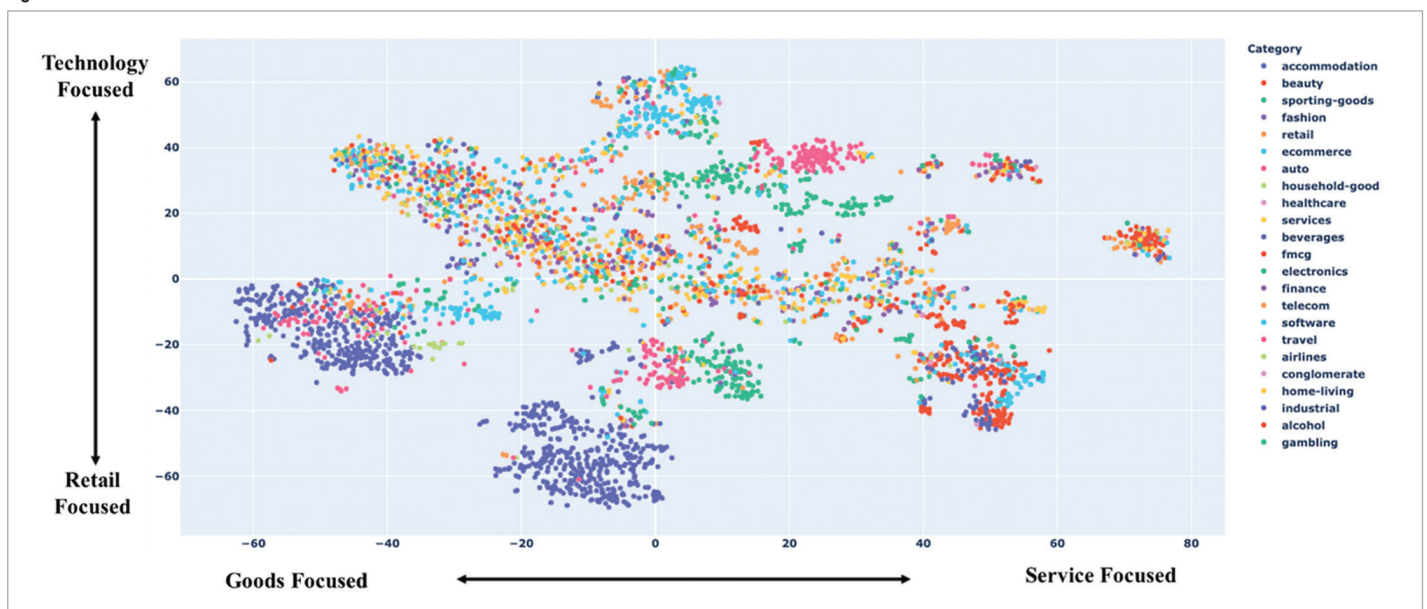
proximities, brand-user proximities, and user proximities as much as possible.[4]

Specifically, in encoding these brands and users, we seek to preserve two network structures: proximity to neighbors and proximity to neighbors of neighbors.

The autoencoder creates a bottleneck through which the input data passes and, using this encoded data as input, constructs a representation of the original data while reducing unuseful complexity. We think of this as the autoencoder training the network to ignore the noise in the data and focus on the primary latent structure in the network data, a method comparable to principal component analysis (PCA) or multi-dimensional scaling (MDS).[5]

For purposes of developing market structure, it is this bottleneck-reduced encoding that interests us. It allows us to identify and visualize the product markets downstream (See Figure 1).

Visualizing the Structure

We next graphed the global structure of the brands in our Facebook social engagement data. Each data point in Figure 2 represents a brand belonging to one of the twenty-five Facebook categories, with each category indicated by a different color. The closer any two brands are in the figure, the more similar their brand representations are. The two dimensions of the map illustrate the contrast between service- to goods-focused categories (x-axis) and retail to technology-focused categories (y-axis).[6] The map reveals some salient clusters of categories.

Zooming in, in Figure 3, we see non-luxury domestic and imported automobile brands such as Toyota, Nissan, and Mazda, as well as automobile accessory brands like General Tire, Auto Alliance, and Auto Parts in one area. Meanwhile several luxury automobile brands including BMW, Mercedes-Benz, Audi,

Figure 3. Automobile Brands and Their Vicinity
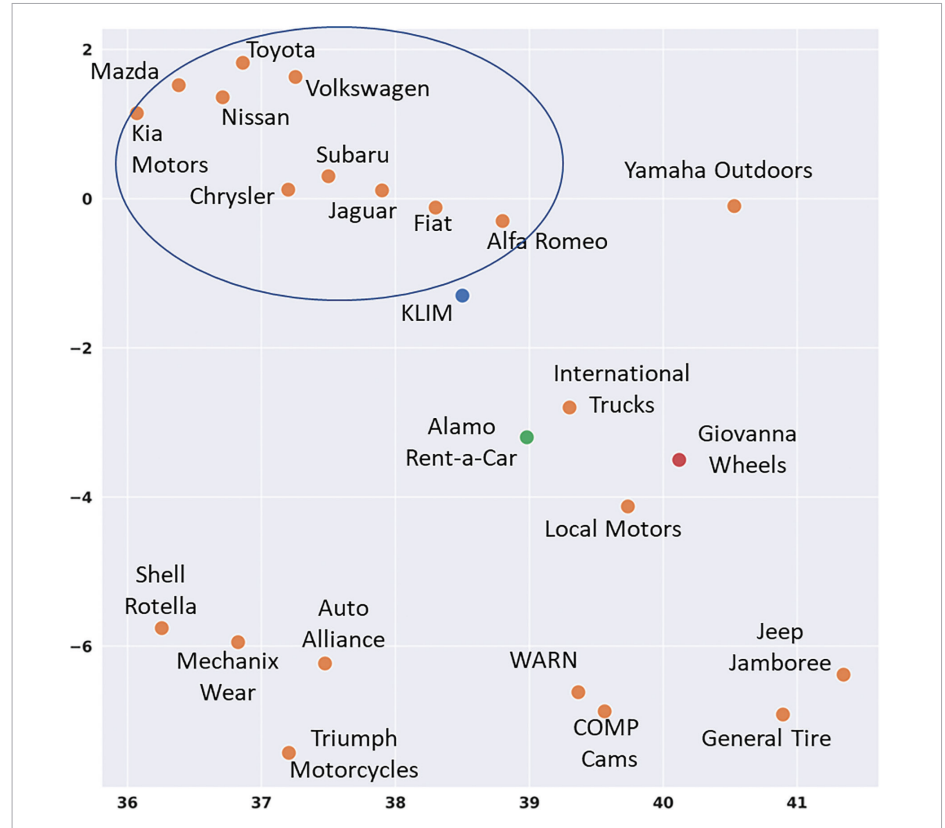
Figure 4. Airline Brands and Their Vicinity

Table 1. Proximal Brands					
Focal brand		**United**	**Southwest Airlines**	**Audi USA**	**Nissan**
Rank	1	American	JetBlue	Mercedes-Benz USA	Mazda
	2	Delta	Frontier	BMW USA	Toyota
	3	Lufthansa	Allegiant	Land Rover	Volkswagen
	4	Southwest	Delta	Lexus	Kia Motors America
	5	Alaska	Alaska	Chevrolet Camaro	Subaru of America
	6	All Nippon	United	Maserati USA	Chrysler
	7	Air China	Airfarewatchdog	Kawasaki USA	FIAT
	8	LATAM	American	Firestone Tires	Jaguar
	9	Air New Zealand	Virgin America	Tesla	Alfa Romeo
	10	Airfare-watchdog	Hyatt	Ram Trucks	KLIM

Tesla, and Maserati are clustered in a different region of the map, with other luxury brands like Channel, Gucci, and Cartier.

This separation between luxury car brands and non-luxury car brands further confirms that our approach captures brand representation in multiple dimensions, not just in industry verticals like product categories, but also in price, luxury, and more. The strength of our methodology lies in how easily it expresses all these relationships on a single map.

Another zoomed area, Figure 4, shows airline brands as well as some hotel and cruise brands that have complementary relationships with airlines. These maps validate our methodology, revealing the core brands that make up an industry as well as the overlaps between markets. Disney Cruise Line, Hyatt, and Southwest Airlines, for example, appear near each other in the circled area, indicating the possibility of joint promotions. This map also shows some customer segments.

Finding Proximal Brands

Visual mapping provides a gestalt of all 5,000-plus brands in the aggregate, but it does not show the distance between the brand vectors in the reduced space. Since identifying proximal brands for substitute or complement analysis is critical to marketing decisions, we identify proximal brands from the perspective of a focal brand. This perspective reflects the various relationships in the social engagement space, from substitute to complement.

In Table 1, we chose United Airlines and Southwest Airlines from the airlines category and Audi USA and Nissan from the automobile category because these brands are generally considered to have different consumer bases and to belong to different sub-markets. Using each of these as a focal brand, we found their top ten proximal brands based on similarity between their vectors.

These proximal brands reveal several interesting points. First, that our method reveals the latent characteristics of specific brands. Southwest Airlines is generally considered a lower-budget airline than United. An examination of proximal brands reflects this difference. The brands closest to Southwest are JetBlue, Frontier Airline, and Allegiant, while those closest to United are major domestic and international airlines, including American, Delta, Lufthansa, All Nippon, Air China, LATAM Airlines, and Air New Zealand. Second, that we can observe asymmetric competition. Southwest

is the fourth closest brand to United, while United is the sixth closest to Southwest.

Third, by analyzing social engagement structure, we discover that brands close to each of our focal brands represent different industries. A brand called Airfarewatchdog, a flight deal finder with over a million followers on Facebook, is close to both United and Southwest Airlines. Traditional market analyses could simply ignore this brand, since it is not an airline.[8] And Southwest Airlines is closer to Airfarewatchdog than United is, which may indicate that fans of Southwest Airlines are more likely to use a deal finder before purchasing flight tickets. Airfarewatchdog could thus be a complement to Southwest directing customers to Southwest's cheap flights, or it could compete with Southwest by directing customers to other airlines.

> Our social engagement mapping can help managers to spot brands outside their product market that are close to a given brand.

Identifying Opportunities and Threats

Our social engagement mapping can help managers to spot brands outside their product market that are close to a given brand, and thus see the opportunities and threats posed by different brands. Consider the airline product market as an example.

Our analysis identifies Disney Cruise Line and Hyatt, both outside the airline market, as proximal to Southwest but not United. This proximity arises from the greater number of users in our dataset who liked both Southwest and Hyatt (2,709) (Segment 2 in Figure 4) compared

with those who liked both United and Hyatt (954). Similarly, a greater number of users liked both Southwest and Disney Cruise Line (3,050) (Segment 1 in Figure 4) than liked both United and Disney Cruise (729).

These proximities reveal opportunities for Southwest to target users who liked Disney Cruise and Hyatt on social media. Southwest could cross-promote with Disney Cruise and perhaps Hyatt on each other's websites. They could also launch coalition loyalty programs to take advantage of their common user base.

From the viewpoint of Hyatt's competitors, these possibilities could be threats, so the same information could help them to take preemptive action. These opportunities and threats are hard to spot using pre-specified categories, so they are difficult, if not impossible, to identify through other means.

When Brand Relationships Change

The structure of markets evolves over time but can change rapidly, especially under an unexpected shock. By learning adaptively from such changes, our method could provide useful insights to practitioners. We analyzed how the structure of markets changes under the influence of outside shocks, using Amazon's acquisition of Whole Foods and Tesla's introduction of the Model 3 as case studies.

We used data from three months before the event was announced and three months after to calculate changes in the distance between the focal brands (Amazon and Tesla) and other brands selected from the same category.

We hoped to discover how a major event changes the focal brand's relationship with other brands. We selected several brands from the retail and e-commerce category for comparison

Figure 5. Changes in Cosine Similarities for Amazon
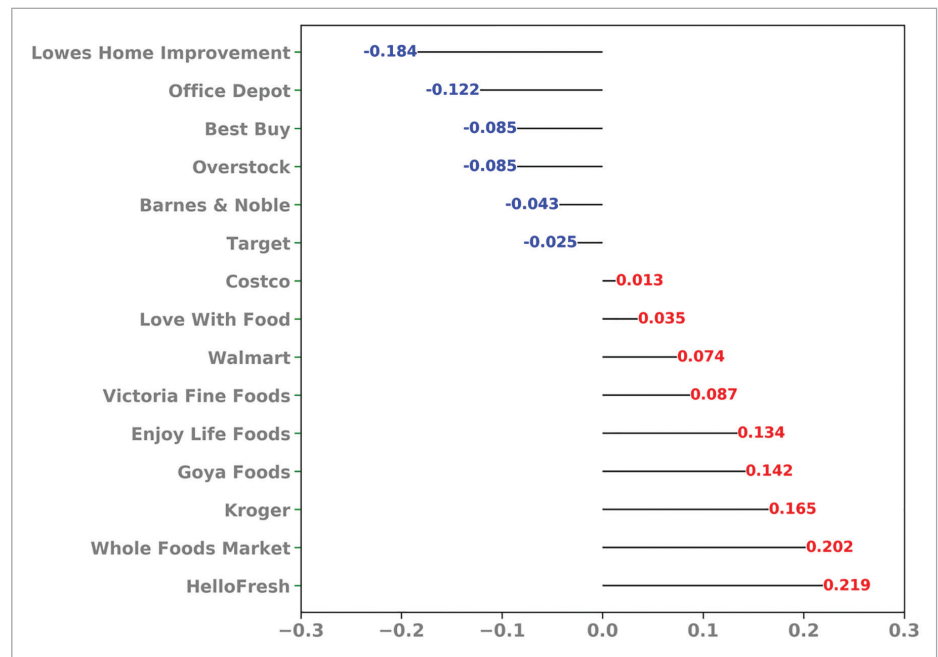
with Amazon-Whole Foods, and several automobile brands for comparison with Tesla.

We investigate the AI's understanding of the change in proximity between focal brand *i* and target brand *j* after the event. We used the cosine similarity[9] of its before and after representations to measure this change: CosSim(after) - CosSim(before). Positive numbers indicate an increase in similarity while negative numbers indicate a decrease.

Amazon Acquires Whole Foods

Amazon acquired Whole Foods in June 2017. The event had a significant impact on grocery and retail industries. At the time, many believed that Amazon's plan for Whole Foods was to fulfill online orders by entering the offline grocery delivery business. (Amazon and Whole Foods ran separate Facebook pages.) After the merger, the system shows that Amazon moved closer to retail brands, while its proximity to other brands decreased slightly.

In Figure 6, for example, the cosine similarity between Amazon and Lowe's Home Improvement de-

creased by 0.184. Meanwhile, the cosine similarity between Amazon and other supermarket retailers increased. The proximity of Amazon to Whole Foods increased by 0.202, and that between Amazon and Kroger by 0.165.

Our model showed that Amazon even moved closer to Walmart, indicating that Amazon's competitive market landscape has shifted. We also found that, after the Whole Foods acquisition, the number of common users who interacted with both Amazon and Whole Foods on their Facebook public pages increased.

In short, after Amazon acquired Whole Foods, online social media users who were Amazon's fans paid more attention to Whole Foods, and users who were fans of other supermarket brands engaged more with Whole Foods as well. The deep autoencoder captures these dynamics and updates the brand representation accordingly.

Being acquired by Amazon had an impact on the market structure of Whole Foods, too. When we examined

Figure 6. Changes in cosine similarities for Tesla

Figure 6. Changes in cosine similarities for Tesla
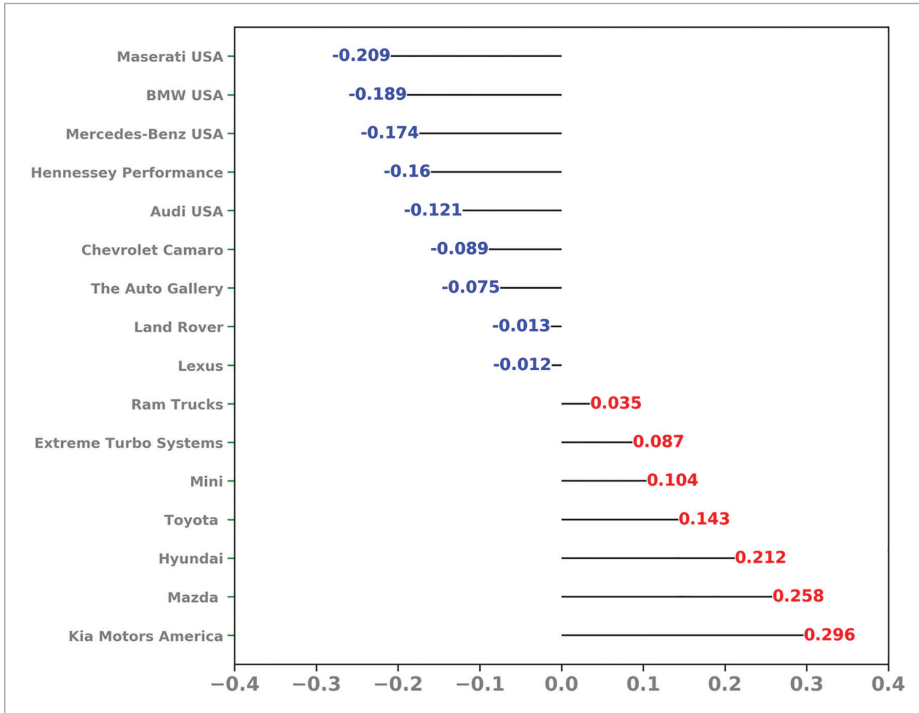
Whole Foods as the focal brand and calculated the change in its proximities to other brands before and after, we found that Whole Foods' proximity to other retail brands such as Target, Walmart, and Best Buy increased.

Perhaps unsurprisingly, Whole Foods' proximity to Amazon rose the most, with the increase in the number of users who liked both. Meanwhile its proximity to supermarket brands such as Goya Foods, Enjoy Life Foods, and HelloFresh decreased slightly.

The magnitude of change in Whole Foods' proximity to other brands was smaller than that of Amazon. This difference seems to indicate that the acquisition affected Whole Foods less, leaving it still positioned near other supermarket brands, while Amazon expanded closer to the grocery retail category.

While this analysis is retrospective, it demonstrates how our approach offers managers a series of snapshots of the structure by which to measure changes in a brand's relative position, allowing them to identify potential shifts in the market structure as social engagement with these brands changes.

Suppose, for example, that the leaders of supermarket chain A observe that Amazon is moving closer to its position on the map. This shift may indicate that Amazon is getting more likes or comments from A's customers. Since one motivation for liking a brand on Facebook is to receive some benefit, like a coupon or discount, it could specifically indicate that Amazon is conducting effective promotional marketing campaigns on social media.

Whatever the underlying reasons, the increasing proximity of Amazon on the brand map can give A's marketing managers early warning of the potential threat.

Tesla Announces the Model 3

Tesla sells two types of sedans, the Model S and the Model 3. The Model S, released first, is a luxury premium sedan with a greater range of acceleration and customization options.

The Model 3 is a more affordable mass-market electric vehicle. The Model S can cost over $100,000 depending on the configuration, while the Model 3 costs about $35,000. Our method reveals that, after unveiling the Model 3, Tesla moved further from luxury car brands and closer to non-luxury car brands.

In Figure 6, we can see in that the cosine similarity between Tesla and the luxury car brand Maserati decreased by 0.209, and that the proximity between Tesla and other high-end or luxury car brands, such as BMW, Mercedes-Benz, Audi, and so forth, changed in similar ways. Meanwhile, Tesla drew closer to Kia, Mazda, and other more affordable car brands.

Potential Uses of the Social Engagement Structure

As brands increasingly use social media to engage, new opportunities arise, allowing them to interact with customers, better understand their preferences, and serve them better. They also see new opportunities to use data describing engagement with the broader market to learn about customers' affinities for brands both inside and outside their own product markets.

In addition to mapping the broader social engagement structure, our methodology can also produce insights that will help managers decide which segments of their customer base to target.

Returning to Figure 4, we can locate user groups on the same maps we use for brands, identifying those most receptive to marketing from nearby brands (Segments 1 and 2). Disney Cruise Line and Hyatt are outside of the airline market but are proximal to Southwest and not United. Identifying clusters of users who are also proximal can help these companies to target them accordingly.

Another important strategic use of our market structure maps

is to identify competitors and complementors across industries and track how these relationships change over time. We provide a more dynamic structure than previous methods, rooted in actual customer or user social media activity.

Our market structure map is also more prescient than earlier methods, predicting emerging competition and complementors by focusing on each stage of the customer's purchase journey, through many categories. By understanding customers' affinities to brands without being confined by product category, managers can

> ## The power of our method lies in its ability to capture the dynamic changes in market structure.

keep alert to impending opportunities and threats.

The power of our method lies in its ability to capture the dynamic changes in market structure. We recommend that firms make such analyses of social engagement

structure, and of the related maps, part of their routine environment and market monitoring and intelligence gathering.

By applying this method on a quarterly basis, managers will be able to spot changes without being distracted by noise in the data. Our method is readily generalized to other platforms, as long as we can construct a heterogeneous brand-user network from the engagement data of public pages. This versatility opens new vistas of opportunity for gathering powerful information from customers' social engagements. ■

Author Bios

P. K. Kannan is the Dean's Chair in Marketing Science and the Associate Dean for Strategic Initiatives at the Robert H. Smith School of Business at the University of Maryland. His research is on marketing modeling, applying statistical, econometric, machine learning, and AI methods to marketing data. He has won several prestigious awards and grants. He is a Distinguished Scholar-Teacher at the University of Maryland and consults widely for nonprofit and profit organizations.

Yi Yang is an Assistant Professor in the Department of Information Systems, Business Statistics and Operations Management (ISOM), School of Business and Management, at the Hong Kong University of Science and Technology (HKUST). His research interests include natural language processing (NLP), machine learning, and statistical inference, as well as their applications in management science. He has received several research grants from the Hong Kong government. He has been consulting for a leading hedge fund firm on financial NLP topics.

Kunpeng Zhang is an Assistant Professor at the Robert H. Smith School of Business at the University of Maryland. His research concerns developing and applying machine learning algorithms to analyze large-scale unstructured data in social media, healthcare, and finance. He has won several research and teaching awards. He has also received several research grants from government and industry agencies.

Endnotes

1. Yang, Y., Zhang, K., and Kannan, P. K. (2022). "Identifying Market Structure: A Deep Network Representation Learning of Social Engagement." Journal of Marketing. 86(4):37-56. doi:10.1177/00222429211033585

2. Hoberg, Gerard, and Gordon Phillips (2010), "Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis," The Review of Financial Studies, 23 (10), 3773–3811.

3. Zhang, Kunpeng, Siddhartha Bhattacharyya, and Sudha Ram (2016), "Large-Scale Network Analysis for Online Social Brand Advertising," MIS Quarterly, 40 (4).

4. The system learns through a deep autoencoder, an unsupervised learning model consisting of two joint components, an encoder and a decoder. The encoder, fueled by a FFNN (fully-connected feed-forward neural network), is a compressor that transforms input data into a compressed low-dimensional latent representation, while the decoder reconstructs it back to the original input data. Because the input data is often high dimensional, such as images or text, in a magnitude of millions, learning effective low dimensional representation (in a magnitude of low hundreds, typically 300) in an efficient way while preserving information in the input data as much as possible is not trivial. See also: Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean (2013), "Efficient Estimation of Word Representations in Vector Space," arXiv:1301.3781 [cs].

5. While PCA reduces dimensions by linear combinations of the input variables, the autoencoder's reduced dimensions are non-linear and non-orthogonal, which are achieved through non-linear activations of the neurons, allowing the model to learn more powerful generalizations than PCA can. We can also compare it to multi-dimensional scaling (MDS), which reduces the information while preserving brand similarity data.
While MDS can deal with data in the thousands, the Autoencoder can handle big data in the millions. Extant literature shows that the Autoencoder performs better than MDS and provides generally different results. For further information about our approach, technical details, and a discussion of the validation of the results, see Yang, Zhang and Kannan (2022).

6. Maaten, Laurens van der, and Geoffrey Hinton (2008), "Visualizing data using t-SNE," Journal of Machine Learning Research, 9 (Nov): 2579–2605.

7. https://market-structure.github.io/index.html

8. Shugan, Steven M. (2014), "Market Structure Research," in The History of Marketing Science, 129–64.

9. https://en.wikipedia.org/wiki/Cosine_similarity

How Artificial Intelligence Can Keep Classical Music in Business

David Serkin Ludwig
The Juilliard School

Many fear that AI spells the end of human music and musicians. David Serkin Ludwig explains why new technologies represent an extraordinary opportunity, particularly for classical musicians, to expand their audience and share the art they love.

Long before I was appointed dean and director of music at The Juilliard School I gave a course in 20th century music history, and for just one class, I would slip in as a listening example an excerpt from Igor Stravinsky's early ballet *The Firebird* – except instead of being a recording of human musicians, it was of high-quality computer synthesized instruments. I would see if any of the students noticed that what they just heard was not an actual orchestra, but

a sophisticated ensemble of digitally rendered performers making sounds that came from zeros and ones rather than 10,000 hours of practice.

Inevitably, most of the students did *not* realize it was a computer playing, and those few that knew something was off explained it only by noting that the recording did not sound "right." This exercise pre-dated many of the newer applications of artificial intelligence in realizing musical scores, but already illustrated how computers can be programmed to perform and ultimately create some very human sounding music.

We have been surrounded by computer-realized music for decades – in supermarkets, movies, and in our vehicles. The cheesy synth sounds of the 60s, 70s, and 80s have given way to sophisticated and convincing instrumental replicants. My composition students are writing music on their laptops using software that would fool the vast majority of people into thinking that what they were listening to was recorded in a live sound studio by living musicians. This realistic playback is assisted by AI algorithms to recreate the acoustics of concert halls and the idiosyncrasies of human performers with astonishing verisimilitude.

But rather than replacing creators and performers, as many fear, AI will be a tool to help musicians write, play, and teach more effectively and engage audiences on a far greater scale. Some of these opportunities have been around for decades but remain largely unrealized, while others are just arriving now. All have implications not only for classical music but for markets where consumers are looking for authentic, high-quality deliverables.

But first, let me clarify the term 'classical music,' which is a terrible descriptor. The so-called 'classical' era in music is also the shortest, from around 1750 to 1820, and two of its most prominent composers were Mozart and Haydn. Today, when we use the term 'classical music,' we are referring to a huge body of work going back 1,200 years, from the first written Medieval chants to pieces written by a living composer yesterday. I prefer the term 'concert music' to suggest music best heard live and best realized in concert rather than in a recording studio. Many musical observers also include jazz as concert music.

When we look into trends in audience engagement with classical music we find much ambiguity. There are studies showing that classical music is primarily and always has been consumed by older generations, but others suggest that millennials and GenZs are increasingly tuning in. One discouraging report says classical music represents only 1 percent of overall music sales, whereas a hopeful one says 30 percent of Americans have *some* connection to the art.[1] Live concert attendance numbers are demonstrably down across the board, but this is true for many cultural events. This ongoing ambiguity has caused classical music institutions to struggle to assert their relevance and value to society for decades.

Threats, opportunities, and robot overlords

With the developments in AI-generated music in mind, the next shareholder-friendly step for commercial record labels would be to cut out the middleman and train computers to make Top 40 music by themselves. This step would greatly lower costs and recording time, but of course would put human creators out of work. While programmable imitations of legendary composers do not usually fool trained musicians, for most listeners a computer simulacrum of Bach sounds pretty much like Bach.

And in popular music, where songs are often created by committees of writers relying on familiar and predictable patterns to make hits, AI can effectively emulate the bands and solo pop artists we hear all around us.[2] Amongst a near limitless number of examples of this on YouTube is a song written seven years ago by AI in the style of the Beatles called "Daddy's Car" (to me, the vocals say middle-to-late period Fab Four).[3] When I have played this track in public lectures and demonstrations, people are

---

The opportunities for using AI-generated music in marketing and advertisement are considerable, as are the questions raised by the practice.

---

astonished to realize that a computer wrote not only the music, but also the words, and performed every note, too. Since "Daddy's Car," there have been huge advances such as Google's MusicLM, which allow a user to simply enter descriptors (text prompts) or even an image, and the AI will compose a new work based on these inputs.

The opportunities for using AI-generated music in marketing and advertisement are considerable, as are the questions raised by the practice. Compensation for their intellectual property is a major concern for musicians because royalties represent an important portion of their income. Who gets paid for AI-composed music that effectively replicates the sound of a band or composer? The developer? The program itself, somehow? These questions come

up every time AI is used for a widespread application in social media. At the time of this writing the controversy is around an app that creates fantasy genre user portraits for Instagram and appears to draw its templates from living human artists who are, yet again, uncompensated.

What rights are left

Legend has it that the composer Victor Herbert founded the performing rights organization ASCAP (American Society of Composers, Authors, and Publishers) after hearing a pianist playing his music all night to entertain patrons of a hotel. Observing that the establishment was making money off his hard work, Herbert started ASCAP to protect his own intellectual property and that of his fellow composers. It is not hard to imagine what he would think of music composed by a computer in his inimitable style to advertise cars, candy, and cat toys, all without a cent going to him.

No one in the arts is immune to the pressing concern that every job is about to be taken away by Skynet,[4] and this concern is both warranted and understandable. We have seen the impact of computer-realized scores in the Hollywood film business, where the number of studio-session musicians in LA has substantially diminished over the past thirty years.

Visuals of human musicians are being simulated as well. There is a cottage industry – already a decade old – of recreating pop stars from the past as projected images so that beloved performing artists can take the stage long after they are gone. You can see convincing holographic replicas of Tupac, Amy Winehouse, Elvis, and Liberace perform in a stadium with a professional band of human musicians backing them up.

Holographic performance is not limited to actual *people*, either. Hatsune Miku is an anime character who dances and sings in massively attended concerts (her voice is an AI-synthesized amalgam of human voices). It was estimated that the brand had generated a revenue of around ¥10 billion in 2012, just five years after the character was 'born.' Miku has since topped charts in Asia, the U.S., and Europe, and was most recently slated to appear at the 2020 Coachella Festival before it was canceled due to COVID (ironic, given that the artist herself is vulnerable only to computer viruses).

Pandora's Bachs

When technology closes one door it usually opens another, and artists are using the Internet and AI to cultivate audiences and market their work with great ingenuity. Since the new millennium, the Internet has widened our engagement into global cohorts, and this shift has changed music industry marketing in ways you might not expect. While music consumers used to be fed what they heard on the radio or listen to new recordings suggested by friends, many people today stumble upon artists new to them through algorithmic recommendations.

YouTube and Instagram are the most important social media platform for classical musicians, and both allow practitioners to not only offer snippets of their performances, but to share their bona fides with the world in the form of personal pictures and their most recent accomplishments: "I am humbled to receive this honor…." Classical music has not yet translated to TikTok as successfully as more popular forms of music, but it is only a matter of time.

One way that classical musicians are engaging audiences online is by posting practice videos that compel social media's AI algorithms to pay attention. The iconic violinist Hilary Hahn is a friend from my student days. Along with her extraordinary musical abilities, she has an enviable savoir faire about engaging her followers as both a performer and teacher.

Hilary invented the "One Hundred Days of Practice" initiative and hashtag to accompany the eponymous goal with daily practice videos. This birthed a movement in which musicians from all walks to set work goals and accomplish them in front of their online followers. There is an authenticity in "One Hundred Days of Practice" that would take significant effort to replicate with AI – and little market incentive to do so.

> AI and VR can help convert followers into audience members.

Best (virtual) seat in the (digital) house

A strong social media presence does not necessarily translate into ticket sales, but AI and virtual reality (VR) can help convert followers into audience members. For years, audiences have been able to tune in to concerts via livestreams from schools and presenting organizations. This was once a revolutionary idea, and many feared its effect on 'butts in seats.' We have since learned that greater engagement of any kind seems to improve attendance for live events, and livestreamed concerts are no longer a rare novelty but ubiquitous content in a saturated market of low production videos.

There is an exciting future in creating virtual concert spaces which are far more engaging than watching a video on your iPad. Imagine your VR headset taking you to a

legendary concert hall far away to hear your favorite musicians perform in a visually believable environment. You are immersed in the sound of a fully recreated space modeled on the acoustics of the hall itself, specifically responsive to the shape of your ears, and with the sound source adjusting to your every movement.

Add some virtual friends to sit with you and you have the beginnings of an unforgettable but remote concert experience, with no one for you to disturb when it is time for a bathroom break. The whole event will be tailored by AI according to your feedback, and if you do not hear or see well, the algorithms and hardware can address that, too.

Perhaps you do not want a seat in the audience for the performance. Perhaps you would like to see what the principal clarinetist is seeing from the orchestra as the conductor elicits that lyrical solo in Rachmaninoff's second symphony. Or maybe, (spoiler alert), while watching Puccini's *La Bohème,* you would like to drop in from the point of view of Rodolfo calling out to his beloved Mimi on her death bed.

AI will help bring you to the space, note your preferences, and provide you with a custom virtual concert experience, though presumably with advertisements for the opera company, suggested recordings, and whatever snacks you ordered on Amazon that week.

I recall a heavy metal concert in World of Warcraft, the popular massively multiplayer online role-playing game (MMORPG), a decade ago. You could send your avatar, be it a warrior, sorcerer, or shaman, to the town square to hear a concert in which the performers transformed into characters themselves. The video game company created a way to engage its 12 million subscribers with a live concert happening in its fantasy universe of castles and dragons. The virtual

concert idea was there, but still existed in a cartoon world.

The technologies of the fully virtual concert hall, with spatial audio, a visually convincing environment, and movable perspective are more or less available now, but no one has yet put in the time and resources to combine everything into the total experience. If someone does build a realistic virtual space for classical music or jazz, it is more or less guaranteed that there is a market audience of eager listeners who will pay to be there.

Compose yourself

My composition students are writing music with the help of AI algorithms, whether they know it or not, by using notation programs and synthesized sound modules for playback. When I was a student twenty years ago, I would bring to my teacher handwritten music on staff paper to be heard in our inner ear or realized on the piano.

Now student composers often write directly into a computer with access to immediate playback with AI instruments through a notation program or digital audio workstation like Ableton or Logic. And you do not have to be a conservatory student to use AI in your composition. For some years the music notation program Sibelius has had an inline ideas panel that suggests variations on new composers' melodies, just as Microsoft's 'Clippy' took us to the thesaurus when we needed a new word.

There is a future where composers work with dedicated AI tools to create unique acoustic worlds for public and personal enjoyment. Just as a chef knows how to mix ingredients, a composer learns to mix sounds. And if you don't want a custom sound environment made for you, you can do it yourself, just like cooking at home.

Imagine putting on some tunes after work, but instead of your favor-

ite Pandora playlist of commercial recordings, you used a simple composition app to create your own music, tailored specifically for those difficult Mondays. Maybe your musical concoction will be in the style of your favorite composer or band, or perhaps you will prefer a more 'spatial' sound environment like a musical analogue to a Rothko painting.

> There is a future where composers work with dedicated AI tools to create unique acoustic worlds for public and personal enjoyment.

Either way, AI-assisted tools can give the consumer agency in making music, and that kind of access to the creative process inspires users to take greater interest in professionally presented music as well. There are currently 74,800 people in the United States who describe themselves as composers of classical music, including film- and videogame-scores.[5] AI composing tools could expand that number by a hundredfold.

I, Teacher

Many, if not most, musicians teach in part to supplement their income. Around the world, young and old alike have a strong interest in learning to play musical instruments. In the United Kingdom alone, over a million adults say they took up an instrument during the COVID crisis. Seventy-four percent of those say that music is important to their quality of life and nearly 40 percent say that this importance increased during lockdown.[6] Many of those new and returning instrumentalists are studying classical music.

The skills required to play classical music professionally are varied and deep, and the ability to teach those skills is a valuable

asset, worthy of cultivation. Musicians who make a living by teaching privately have a limited number of hours a week to devote to students, But what if an AI teaching partner could expand that capacity? An application could be programmed to observe a student's progress and offer feedback, mapping body movements and embouchures.

Much as a nurse prepares patient information for a doctor for more efficient diagnosis, an AI assistant could convey the metrics of a practice session to a music instructor. A programmable teaching partner could also double as a practice buddy, keeping the student company during the tedious work of playing scale exercises, not just notifying the student when they are sharp or flat, but using the pre-programmed guidance of a human instructor to tell them how to fix it.

This virtual teaching partner could function far beyond the meters and gauges we use now to show when we are playing in tune. We could communicate directly with an AI anime character or computer-generated human who would help us to play better. And we could use virtual students to train teachers, seeing how our pedagogical methods work in order to instill good instruction habits.

Not just for technical instruction, AI can also act as a collaborative performing partner. Classical musicians have tried, mostly unsuccessfully, to play together over an Internet limited by bandwidth and equipment. Meanwhile conservatories scrambled to find performing opportunities for students trapped at home during the pandemic.

> We live in a world where an aspiring musician can develop their performance chops with a robot duettist who lives in their phone.

Computer accompaniment technologies that have existed for decades have evolved to be more responsive and expressive companions. There are many examples of AI-accompaniment software. Among them is Juho Pohjonen's MyPianist app that listens to a musician's phrasing, volume, and tempo and responds musically by changing its own tempo or volume just as a human collaborator would. We live in a world where an aspiring musician can develop their performance chops with a robot duettist who lives in their phone.

Ahead of the beat

The innovations of living musicians are essential to performance, composition, and instruction, but can musical creativity itself be modeled? That we even have to ask suggests that the answer is yes. But the intangible combination of intuition and order that goes into human expressive genius will not be replaced any time soon. Maybe we can simulate Beethoven well enough to convince most of the people most of the time, but some will always seek authentic performances and have the acumen to recognize the real thing when they hear it. These sensitive listeners and the pleasure they feel in experiencing live concert music is the brass ring of 21st century careers for classical and jazz musicians.

Classical music has no choice but to embrace technology and make the most out of it. In an era of instant gratification, this vast body of work, which requires dedicated study to play professionally and at least some exposure to connect meaningfully with, could fade away as a literature, leaving only the well-known hits behind. Eventually we would be left with nothing but the opening movement of 'Eine Kleine Nachtmusik.'[7]

Classical music is often attacked, perhaps most often by classical musicians, as being too elitist. But what happens when what is considered elitist in art is made accessible to everybody? The greatest hope for this music's future is for its practitioners to learn to wield AI to engage audiences and create inclusive ways for a wide variety of listeners to enjoy their work. Where some fear the ultimate death of art itself in the growth of artificial intelligence, I see an extraordinary opportunity to share the art we are so passionate about with a global audience. And a desire to share what we love is, after all, only human. ◼

Author Bio

David Serkin Ludwig is the Dean and Director of Music at the Juilliard School in New York City and an award-winning composer whose works have been featured on stages and screens around the world, from Carnegie Hall to a presidential inauguration. Ludwig holds positions with nearly two dozen orchestras and music festivals in the US and abroad and resides in New York City.

Endnotes

1. https://medium.com/@AmericanPublicU/popular-classical-music-how-popular-is-classical-music-part-ii-4040456752db
2. https://www.theguardian.com/music/2012/jul/27/pop-music-sounds-same-survey-reveals and https://www.wpr.org/pop-music-getting-more-repetitive-computer-scientist-says
3. https://www.youtube.com/watch?v=LSHZ_b05W7o
4. https://en.wikipedia.org/wiki/Skynet_(Terminator)
5. https://www.careerexplorer.com/careers/composer/job-market/
6. https://www.classical-music.uk/news/article/one-million-adults-took-up-an-instrument-during-lockdown-uk-music-survey-reveals
7. https://www.youtube.com/watch?v=jLcmwY4hjlw

INFOSYS:

AI Helps Build Customer Engagement to Ace the Tennis Game

Navin Rammohan and Mukul Pandya describe Infosys Courtvision, which allows tennis fans to visualize the data points of their favorite players during the match, seconds after the point is played out. This AI and 3D tech-powered feature has existed since the late 2000s on broadcast television, but never before have fans had access.

Navin Rammohan

Infosys

Mukul Pandya

The Wharton School,
University of Pennsylvania

In 2015, Infosys, an IT services company headquartered in Bangalore, India, was looking for a global sport with which it could get involved to show how innovative technology could drive customer engagement. The company had three objectives. First, it did not want to merely slap its logo or brand on a sports event; instead, it wanted to establish a genuine partnership that could showcase a variety of digital innovations. Second, Infosys hoped to reimagine the sport through data, insights, and digital experiences for all stakeholders - players, coaches, fans, and the media. And third, the company sought a sport that was popular all over the world, and especially in its key markets.
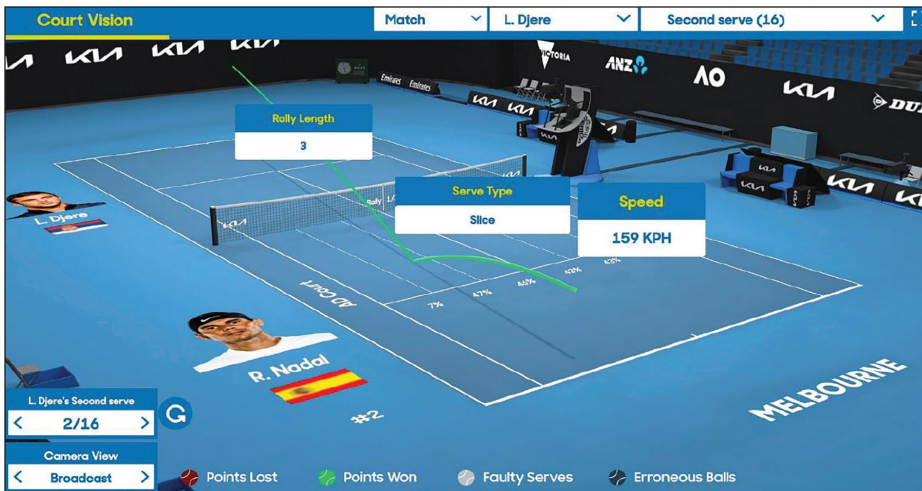
Features like Courtvision have seen close to a million fan engagements on the websites and mobile apps of ATP, Australian Open, and Roland Garros.
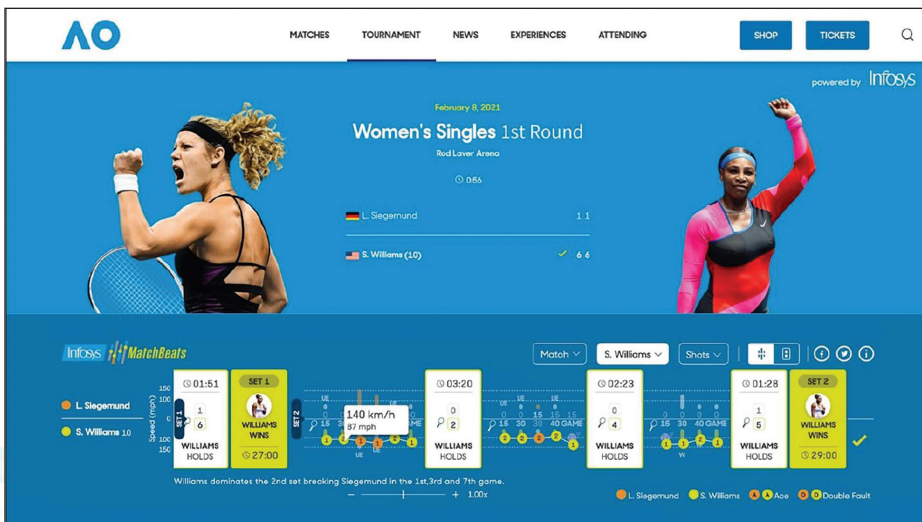
This is an example of the popular feature Infosys MatchBeats as viewed during a live match at the Australian Open in 2021.

After considering various options, including soccer and car racing, Infosys settled on tennis as a global sport that met all these criteria. It is the world's fourth most popular sport, with more than one billion fans. Moreover, many of those fans were in the company's key markets in the Americas, Europe, and the Asia-Pacific region. Infosys believed that, if it got deeply involved with tennis, it would reap many digital transformation opportunities.

In order to achieve these objectives, Infosys collaborated with a diversity of groups and organiza-

tions in the tennis ecosystem. In all, over the past seven years, the company has introduced more than twenty innovations. One example is

Infosys hoped to reimagine the sport through data, insights, and digital experiences for all stakeholders — players, coaches, fans, and the media.

Infosys MatchBeats, which powers live stats from using a foundation of machine learning and automated natural language insights. This feature alone has had more than 75 million fan views since its launch in 2019, and numerous fans have tweeted asking why MatchBeats is not available for all tournaments. (So far, Infosys has implemented MatchBeats at more than sixty tournaments.)

Few tools among tennis coverage innovations have had a particularly impressive impact on customer engagement.

AI Highlights and AI Shot of the Day

The Infosys AI Highlights and AI Shot of the Day features can curate hours of tennis match data, automatically selecting highlights to be made public minutes after the game is over. They also overlay these highlights with graphics that showcase the key moments of the match. Given the speed at which the AI works, it can actually provide these services while the match is underway or after every set or game, as well as serving social media snippets. This immediate response would be impossible without AI.

Using AI, Infosys can look at every facet of the games from the first day of the tournament to the last. It allows the company to pick out the best shot of the tournament, using multiple parameters including the roar of the crowd, the length of a rally, the spin and speed of the ball, and many other factors. If the media teams at the Australian Open, for example, want to look at just aces or forehand shots, the AI can help them do that. The technology can also rapidly create playlists based on these attributes.

To give you an idea of the scale of these features, on Day 5 of the Australian Open 2021 alone, 2,703 points were played in the singles

competition, 1,633 in men's singles and 1,070 in women's singles. Infosys AI analyzed all this data in seconds, using more than 120 filters and over 1,000 combinations to generate suggestions for Shots of the Day, ready-formatted for instant publication on social media such as Twitter, YouTube, Facebook, and Instagram. Innovations such as these have made tennis more appealing to its fans. At the Australian Open 2022, just the content generated by this AI garnered over 10 million media views of the shots of the day.

In addition to these features, Infosys was able to use coding to add an 'elegance value' rating for each shot. Elegance value, or a shot's refinement and style, added subjective, AI-based analysis to the image. That is not to say that a more elegant shot necessarily draws greater customer engagement or vice versa. Artificial intelligence and machine learning (AI/ML) use video analysis, audio analysis, hawk-eye procedural analysis, or a combination of these methods to calculate these ratings in cases in which a shot's elegance is predominantly gauged by the type of shot. Within a single type of shot, they base the rating on the outcome and intrinsic parameters of the shot including speed, spin, proximity to the net, accuracy, crowd noise, and more.

Elegance is a subjective parameter that allows the system to highlight the best shot of every tournament. This concept did not previously exist in tennis; it was made possible by AI and was completely new to the game.

These innovations helped Infosys make tennis coverage more appealing to the media. For example, in collaborating with two of its partner organizations in tennis, Infosys made AI Highlights available on its portal so that every broadcaster around the

The AI Shot of the Day tool's snapshot above has an algorithm parsing the thousands of points played each day to generate the top five strokes in a general category and numerous others. This tool simplifies the lives of media editors and drives engagement, as media teams use it to select shots of the day to post on social media.

Here is an example of an AI auto-generated clip from the Australian Open 2022 that *garnere*d more than *17,000 views*. Similar to this, every day from qualifications to the finals of the tournament, a Shot of the Day was generated and posted on Twitter, Facebook and YouTube.

'HUNGRY' COCO CRUISES INTO MAIDEN SLAM QUARTER-FINAL

MONDAY 7 JUNE 2021 · CHRIS ODDO

Just three years ago, Gauff was in Paris as a junior, where she became the youngest player to claim the girls' singles title since Martina Hingis in 1993.

Since then she has soaked up experience like a sponge. Just two players - Ashleigh Barty and Aryna Sabalenka - have earned more tour-level victories than Gauff in 2021, but neither of them is still standing in the draw in Paris.

"I just feel like I'm getting better and better with each match, and that's kind of how I want it to happen," Gauff said after her third-round victory over Jennifer Brady. "I think that you can tell that I'm improving and making smarter decisions on the court."

Here, Infosys AI-Assisted Journalism let the writer of this report post the story of Coco's dominance on the first serve *during the match*. It turned out to be a critical insight on how she won the match. AI not only generated the infographic and the stats, it also directed the journalist's attention to this influential stat.
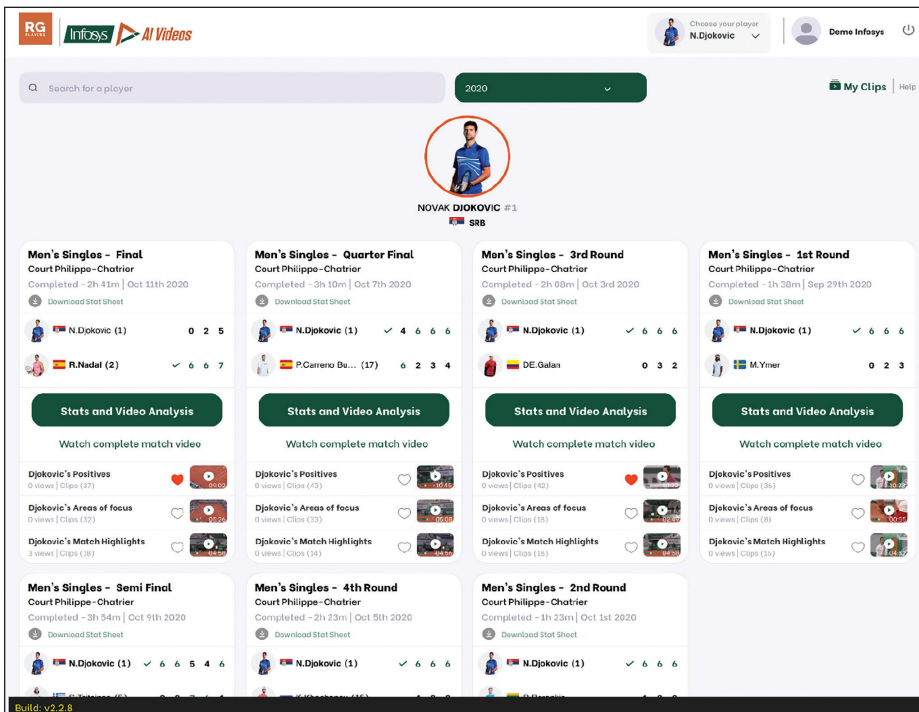
This is a tournament view for tennis great Novak Djokovic and his coaching staff, showing his *positives, areas of focus,* and *highlights* from each match played at Roland-Garros 2022. The AI videos platform for players and coaches has been used by 3,000+ players and coaches. Clicking on "Stats and Video Analysis" leads users to a range of video and stats analysis tools.

world could easily pick it up and instantly embed highlights on digital channels for their audiences. In addition, during one of the Australian Open editions, an organization used AI Highlights to automatically produce instant match highlights minutes after the matches ended

and publish them on popular search engines. This use benefits sports journalists and also offers enormous value to tennis fans everywhere.

Short Video Snippets

Another pragmatic use of these AI video portals is to distribute resulting clips on social media. Short snippets engage fans because they are able to see them instantly on social media. This is another way in which AI has significantly increased fan engagement.

AI-Assisted Journalism

Another innovation that has had tremendous impact is the AI-assisted journalism tool that Infosys created to help journalists instantly create infographics from the match. If users are looking for a certain data point, Infosys can provide it and create an infographic within minutes. This feat was simply not possible in past. Now some journalists use this module to insert infographics into their articles. It also helps to engage fans. For example, at Roland-Garros, also called the French Open, users used this tool to create and post more than 700 match reports. At the 2022 edition of the slam, it drove over 27.5 million interactions on rolandgarros.com.

Did the deployment of AI help Infosys increase customers' engagement with tennis and, indirectly, with the company as well? Absolutely. Before Infosys introduced these innovations, about a million people were engaging with the content. Thanks to these AI tools, the number of fans engaging with the content in 2018 increased by 27 percent from the year before. With the launch of the new match center for The Association of Tennis Professionals (ATP) in 2021, AI has powered a variety of insights and stats-driven features for fans, drawing 1.6 million users.

Overall, 150 million fans have engaged with Infosys and its growth rate is in the high double-digits. Its use increased 38.1 percent in 2022 from the prior year's numbers during Roland-Garros. This increase includes more viewers on tennis websites as well as social media channels. And the key to this sharp increase is that these tools made it easy for fans to consume content.

Player Portal

Along with fans and journalists, players and coaches also benefit from Infosys. The company has created a player portal to help players dynamically view and analyze portions of their matches. In the past, they would have gone to YouTube to watch their entire game performance against a competitor.

With the player portal, they can examine their performance at a granular level. The system allows them to choose exactly what they want to view: for example, all their forehand shots against an particular opponent. This use of AI helps players and coaches to better understand their game, offering subjective information about positives and areas to improve, basic and advanced stats, as well as stats powered by pattern analysis.

AI provides coaches and players with performance data as well as analysis of that data in a way that was not possible before. It is no surprise, then, that more than 3,000 players and coaches have used this tool since its launch just a few years ago, even though it is available only at select tournaments such as the Australian Open and Roland-Garros. From 2021 to 2022, its usage doubled.

AI is also democratizing the use of digital technology.
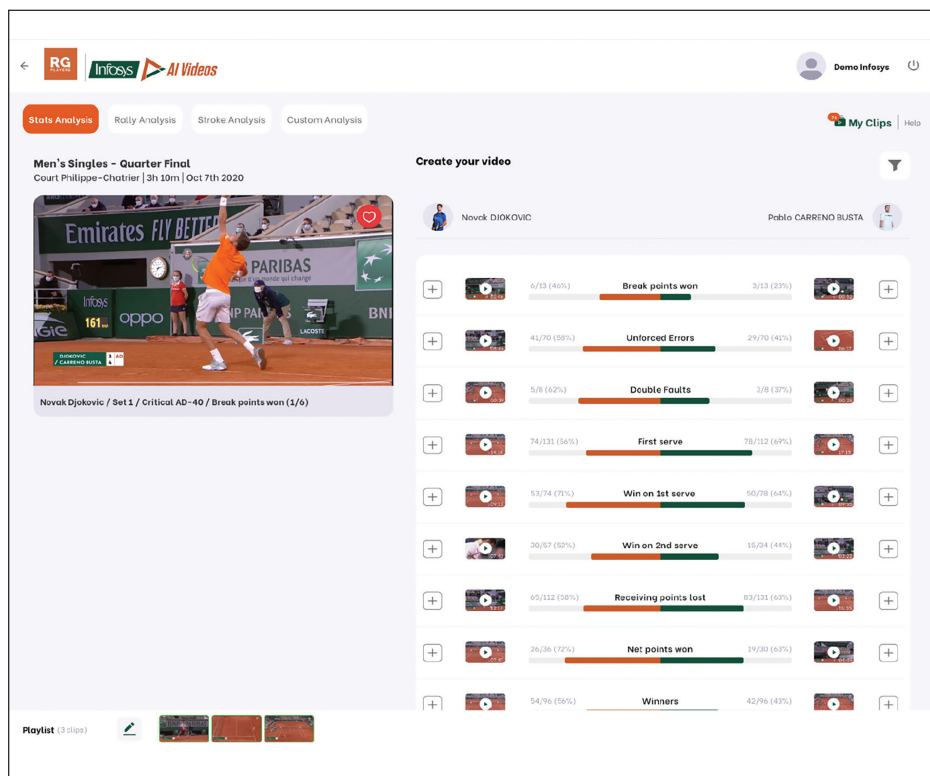
AI is also democratizing the use of digital technology in devising player strategy. In the past, players who could afford data coaches had an edge. Today, technology levels the playing field in data and video analysis, whether the player is ranked 300th in the world or third. This capability grants the sport of tennis some of the most technologically advanced coaching in the world.

Looking ahead

Infosys has increased the engagement of fans, journalists, players, and coaches with tennis. And it seeks to go further. In the future, Infosys would like to personalize the content it distributes, ensuring that each constituent gets the information they most value. At the moment, casual and deeply involved fans get the same content. Infosys hopes to develop an AI model that

can provide exactly the content depth that suits each fan.

Infosys is already making inroads. When the company redesigned the Australian Open app, personalization was at the core of its efforts. For example, it allowed fans to bookmark favorite players and receive customized content about their scores and schedules. Meanwhile, news and push notifications make it convenient to follow the sport.

During the Australian Open 2021, Infosys developed the virtual hub, a digital hospitality platform that offered premium personalized content to key partners of the tournament. It featured virtual rooms with customized welcome messages, partner branding, and unique and differentiated content for various groups. The platform hosted more than 900 senior industry leaders from more than fifty different brands across the globe and delivered more than 19,000 sessions and more than 120 pieces of original content.

Infosys will continue to deliver customized, AI-powered experiences for players, fans, and the media. As the power of AI increases – which users have already seen in the performance of ChatGPT – these tools will only become more intelligent, productive, and usable.

In the future, we will need to focus on creating more sustainable, accessible, AI-powered, and consumer-driven experiences for the tennis ecosystem. Digital sustainability tools will move the company toward a circular economy and tennis toward net zero. Just recently, Infosys launched a new carbon tracker tool to help players achieve net zero, seeking, as always, to improve the experience of everyone who loves the sport. As long as Infosys stays true to this calling, it will continue to ace the AI tennis game. ■

Author Bios

Navin Rammohan is Vice President and Segment Head of Marketing at Infosys, responsible for the company's European marketing initiatives. They also manage Infosys' sponsorship and events activities in Europe. Navin has over twenty years of experience in marketing and communications, with a focus on technology companies and brands. They earned a B.Tech from the Government Engineering College, Thrissur, a post-graduate degree from MICA | The School of Ideas, and completed an executive education program at Stanford University.

Mukul Pandya was the founding editor in chief of Knowledge@Wharton from 1998 until his retirement in 2020. After retiring from K@W, Pandya was a senior fellow with Wharton Customer Analytics and AI for Business. A four-time award winner for investigative journalism, Pandya has published articles in *The New York Times, The Wall Street Journal, The Economist, Time, The Philadelphia Inquirer,* and more. He has written or coauthored four books.

AI FOR CUSTOMER ENGAGEMENT AT GOOGLE

Google's Customer Engagement Leadership Team:

Tim Frank	Qiushuang Zhang
Aastha Gaur	Ravi Narasimhan
Abheek Gupta	Roman Karachinsky
Doris Neubauer	Sandeep Beri
Ian Suttle	Shashi Upadhyay
Leo Cheng	Tony Li
Natalie Mason	Vicky Ge

Google had the opportunity, directly or through its representatives, to apply the best research on artificial intelligence and machine learning to its interactions with customers, transforming its approach and creating more value for its customers. Tim Frank and Google's customer engagement leadership team describe this AI transformation, rooted in prioritizing customers.

Google was founded "to organize the world's information and make it universally accessible and useful." The company is guided by a simple tenet: focus on the user and all else will follow. Today, the company is home to a range of consumer products with more than 1 billion users and business products with more than 1 million users. These products have succeeded by focusing on the user. Each of them has

Figure 1: Customer Success Flywheel

experiences. This new application of AI caused our marketing, sales, and support teams to grow, while we delivered great experiences to our users.

The Customer-First Approach

Customer expectations grow continuously and younger customers drive trends that are widely adopted by all ages. Today, we work to create content with nothing to snag or delay our customers, termed 'frictionless.' Users want all their feeds to be hyper-personalized and their content to be 'snackable,' quick and easy to digest. Few have the time or patience to wait to talk on the phone. Users feel that their digital experiences should just work, and when they do not, should be solved on their terms with little friction or effort. Our customers want one personalized, proactive, and contextually appropriate Google experience wherever they come in contact with Google products, yet in a way that respects privacy. To fulfill that wish we used sophisticated software, building a simple customer centered model to guide how we use AI to engage customers.

The customer success flywheel (Figure 1) shows how we build solutions informed by deep customer understanding, insisting that we provide our customers both value and high satisfaction (CSAT[1]). This process starts with a customer's earliest interaction with Google (middle of the flywheel). As we learn more about our customers' preferences in products, marketing, sales, and support, we can bring them more value. This process creates a positive feedback loop and makes every point of contact or surface customer-aware. We deliver 'one size fits one' solutions globally to every one of our customers. Our success metrics also focus on customers, understanding them in order to bring them value and satisfaction.

also gone through an exponential growth phase which required careful management. We have learned that you cannot solve exponential problems with linear solutions.

The customer engagement organization at Google is the human face of the company, facilitating billions of interactions each month. As the business grows, the number of its customer interactions grows even faster, ranging from a customer filing a ticket online to an in-person meeting with a business client's representative. Our interactions start with tens of thousands of Google representatives providing customer support and extend to custom tools and processes that improve our human interactions at scale. Today, we have taken this one step further, bringing even more value to our customers through AI-enhanced

Figure 2: Framework for Transforming Customer Engagement

Prioritization Framework

We prioritize bringing our customers the best experiences and outcomes, regardless of how we deliver it. One way for us to confirm that we are optimizing properly is to ask if we would be comfortable sharing our metrics dashboards with our customers. As organizations are pressured to focus on metrics within their control, it is easy to optimize for success that is not customer-informed or for the customer's benefit. By anchoring our work in customer understanding, value, and satisfaction, we are able to orient our efforts toward our customers' needs.

To help ourselves think about these opportunities, we built a framework (Figure 2) for transforming customer engagement that showcases the need for AI if we are to break out of the 'incremental' quadrant. It also guides us to think about pushing the limits of what is possible for humans.

The y-axis describes customer value and shows the inherent current limit of human return on investment (ROI), given that labor is not free. This scale reveals that we need to raise productivity, using tools like smart reminders, meeting scheduling, auto-pitch deck creation, and more. The x-axis shows how challenging the necessary tasks are for humans. We need the powerful assistance of AI, especially for tasks whose volume, complexity, and optimization are too great for humans to grasp or undertake.
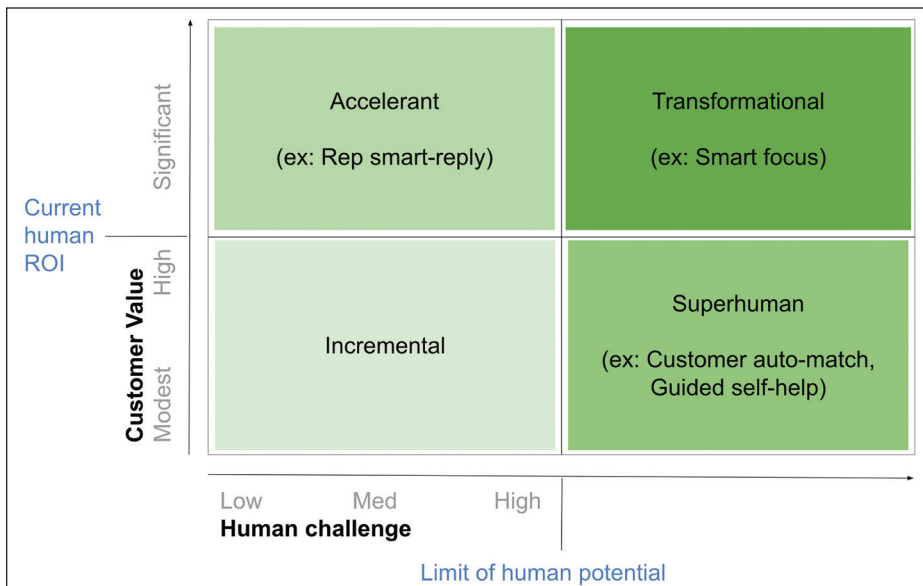
The path to transformation is through foundational AI research. At Alphabet, we are lucky to have world-class research teams like DeepMind[2] and Google Research[3] pioneering AI breakthroughs. We are pushing toward a future in which we can maximize the best of both human and technology intelligence, delivering the most customer value in the long term.

Smart Focus: Engages Customers to Significantly Improve Outcomes

We are developing tools to help our sales and support representatives serve customer needs, hit key performance indicators (KPIs), and help to run a high-performance organization. They have access to insights that can help our customers meet their own business objectives, but because the breadth of knowledge is so great, it can be paralyzing for them to decide where to spend their time. We needed to build AI that would help them to better find the right focus, improving the outcomes of our customers.

An innovative AI application

Past efforts to direct human attention often used hard-coded conditional logic and heuristics to produce average performance; we set out to achieve three things:

1. We used AI to build an intelligent detector, which automatically attunes its sensitivity to different representatives and customers. As a result, its ratio of useful information to filler is very high, so our representatives know that listening will help. Standard anomaly detection would not have worked for such a large variety of customer data on different scales.
2. We converted heuristic-explanation problems into multi-dimensional, multi-metric correlations and attribution models, which significantly reduced the computational complexity. So no matter the suggestion, we can explain the rationale to our representatives, who can then evaluate the logic and explain it to customers or colleagues.
3. We designed our experiments so we could measure the cause of the model's effects and use feedback to train and improve the model.

Resulting customer engagement

These initiatives allow our representatives to start their day knowing whether there is a high priority item ahead. When there is, they can investigate the context, prepare the right response for the specific customer, and be well prepared to help when they converse with the customer.

So far, this tool has improved the outcomes of both customers and representatives. Customers can now avoid unplanned campaign delays and optimize high ROI campaigns while getting more for what they spend on Google.

In aggregate, customers found 23.5 percent more value with Google and 15 to 32 percent greater confidence. We measured this increase by how much customers spent over the subsequent three weeks compared to a randomized controlled trial. Representatives also appreciated the help from the system, reporting 90 percent satisfaction, double the level they reported with non-AI attempts.

Guided Self-Help: Diagnoses and Resolves Issues Using Scalable, Intelligent Support

Each year, billions of customers interact with hundreds of Google's consumer products including search, Android, and YouTube. As Google's offerings and customer base grow more diverse, so does the complexity of identifying and resolving issues.

Customers can ask for help in many different ways, and the problems they encounter may have any number of root causes. One way that Google Support increases self-help offerings, improves representative productivity, and boosts customer satisfaction and resolution is through large language models (LLM).[4]

These promising applications allow us to offer AI-guided support. Our customers can interact

Table 1: Four case studies in which we have made significant progress:				
Case study	Quadrant	Description	Maturity	Customer Engagement and Impact
Smart focus	Transformational	Engages customers when there is significant performance change.	Generally Available	23 percent higher customer value, 90 percent representative customer satisfaction (CSAT)
Guided self-help	Superhuman	Solves customer issues in-product with guided support.	Generally Available	Reduces customer contact rate by 5.3pp, CSAT equal to human help
Customer auto-match	Superhuman	Serves customers faster, shortening the time between signup and value.	Generally Available	Reaches customer value 10x faster, 150 representatives worth of toil eliminated.
Representative smart-reply	Accelerant	Improves the quality of customer service chats.	Generally Available	44 percent decrease in human agent chat, 48 percent increase in customer self help, 95 percent representative CSAT

naturally with intelligent, automated support, such that they feel empowered by and prefer self-help over other support alternatives. We have invested in the two critical steps to addressing customers' problems: diagnosis and resolution.

To diagnose, we find the root cause of an issue by examining its symptoms. To resolve the issue, we use this diagnosis to help our customers choose the best course. This self-help advice is informed by a wealth of information, encompassing help articles and representative transcripts. But customers still bear a significant burden; they must articulate their issue and then navigate a tree of symptoms and causes, trying out different solutions.

Innovative AI application
Our strategy lightens the customer's burden in diagnosis and resolution. For diagnosis, we use cutting-edge LLM applications to efficiently understand the issues as the customer describes them. For resolution, we use a multitask unified model (MUM) to rank and serve solutions through support interfaces like our search page and escalation forms.

These AI applications spare our customers unnecessary complexity. In some cases, we can now take customers directly to the right solution. We have seen a 20 to 40 percent increase in our ability to answer questions depending on product area, as well as a quality of neutral to positive in the answers at all points of customer contact.

Resulting customer engagement
One primary metric we use to measure the success of this self-help tool is the rate of support escalation, instances in which a customer cannot solve the problem with the tool's support and moves on to a human representative. Our experiments showed a 5.3pp reduction[5] in escalations, or 2.8 times the performance of the latest neural networks that use long short-term memory (LSTM).[6] We achieved this contact rate reduction while maintaining our high customer CSAT.

Customer Auto-Match: Serves Customers Faster, Shortening the Time Between Signup and Value

Helping Google Ads customers starts with understanding which campaigns and accounts they are using to achieve their digital marketing goals. At the scale of Google, this is a daunting task. Each year, customers create millions of new Google Ads accounts, and we collect little data about them to avoid adding unnecessary friction. Customers use Google Ads through different legal entities, divisions, third party agencies, and consultants — and there is often no easy way to determine with which person or group Google should engage.

To identify which accounts belonged to each customer, we initially asked several hundred representatives to manually examine each new Ads account and connect it to the right entity in our customer hierarchy. Technically this method worked, but it was slow and difficult to scale to efficiently support the ever-growing set of Ads customers. We needed an automated solution.

Innovative AI application
We decided to implement a set of approaches that would more rapidly match our customers to their accounts:

1. A combination of AI models that would identify the most

likely owner of a new account in our customer hierarchy

2. A rule-based system that allowed our teams to manually require high-confidence rules and signals in handling of simple cases

3. A small team of analysts to identify the most complex and ambiguous accounts and build datasets for machine learning (ML) training and validation

Resulting customer engagement

Customers benefit the most with a 10-fold reduction in the time between signing up and receiving value. Google representatives also appreciate being relieved of some tedious toil. And we have increased productivity by the equivalent of 150+ employees annually.

Representative Smart-Reply: Improves the Quality of Customer Service Chats

Google Support offers personalized one-on-one support over chat, email, and phone to certain consumer and business customers. Initially, we used human representatives but as our range of products becomes more complex, it can be hard for any one person to remember the latest policies for every market and product. We also wanted to expedite resolutions and increase customer satisfaction amidst this growing complexity.

Innovative AI application

We first invested in AI-recommended replies for representatives. During a live chat, the representative sees a suggestion for the next best response to send to the customer. The representative chooses to either approve the AI suggestion or override it with a human input. The system is trained on representative transcripts and optimized by tracking representative acceptance and resolution. Each interaction improves the model's future recommendations.

Figure 3: Technology Architecture for AI in Customer Engagement (Simplified)
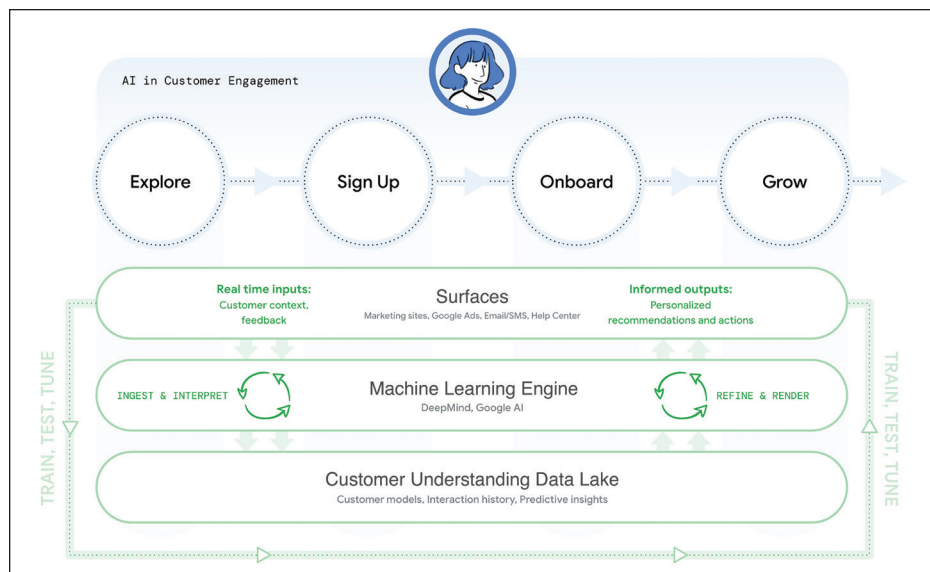
After success with early iterations of this tool, we started conducting supervised experiments for Smart Replies. In this phase, the AI automatically sends its chosen reply to the customer after a certain time, so long as the human representative does not opt out.

This evolution from representative opt-in to opt-out suggests a future in which human representatives can move from handling cases manually to supervising AI with little direct involvement. Human agents could then focus on managing complex customer escalations. Recently, we even observed, for the first time, the AI model handling a complex support issue with no intervention from the human supervisor.

Resulting customer engagement

Google representatives loved the assistance with routine cases. Customer satisfaction reached 95 percent and early results indicated that customers resolved their issues both faster and better. This improvement was indicated by a 44 percent reduction in manual representative messages in chats.[7] The use of smart-reply also led to a 48 percent[8] increase in customer selection of self-guided support

pages. We have therefore reduced our dependence on fallible human efforts in high-pressure, repetitive situations while improving the help we offer our customers.

Google's Technology Architecture for AI in Customer Engagement

To power the case studies above, we built a technology architecture (figure 3) oriented toward the customer's lifecycle, increasingly engaging with customers to better understand their goals while providing assistance and guidance that's easy to use. This architecture has three key components:

1. The Customer Understanding Data Lake collects, curates, and, without sacrificing privacy, scrutinizes data to create a complete picture of our Ads customers and power our AI layer.

2. Our machine learning engine is composed of Google's AI/ML technologies, finely tuned to solve complex and unique problems.

3. Finally, artifical intelligence insights and actions are carefully integrated into various

customer touchpoints, such as marketing sites, help centers, core products, and sales tools, as well as into customer's journeys, actions, and context.

Each element of our architecture was custom-built and optimized to help our customers achieve their business goals with as little friction as possible.

What Leaders Can Learn
Measurement rigor
To confidently invest in a portfolio of AI projects, we must measure the impact of each project and understand which bets paid off. We scrutinize the total impact on our customers and our company, which is a product of both the scale and value of each project.

We run randomized controlled trials (RCTs, or A/B tests) using proven data science techniques, accepting only statistically significant results within a narrow confidence interval, typically 95 percent.

To ensure consistency, we do not allow teams to grade their own work but rather have a centralized, independent group of data scientists and business experts who certify measurements and provide guidance. We are confident that we are delivering clear and quantified value to customers and powering growth in Google's Ads business.

Because of the high rigor required to show statistically relevant results, this indicator inherently lags. To make decisions along the way, we consider other leading indicators such as sentiment, repeat usage, task success, and assistance feedback from both the customer and the representative.

Sentiment is measured both within a workflow (transactional) and outside (relationship). Our AI-powered experiences score 10pp to 20pp higher than non-AI-powered ones. They have been key to

our cracking the 75 percent CSAT threshold.

Repeat usage is an improved measure of coverage and adoption that allows us to detect early promise in a new AI feature. In the smart-reply example, this might be the difference between something being used a few times a day and being used several times in every chat. As we saw in the case of manual chat improvements, performance can be improved by more than 40pp.

Task success is a broad term for achieving a customer's or representative's objective. In sales, it might describe a pitch rate or win rate, while in support it could describe resolution on the first attempt or total resolution time. Historical baselines give us a clear threshold to shoot for and we have seen individual AI solutions increase success rates by ten times on hard-to-improve tasks.

Assistance feedback helps us to understand when an AI failed and required human intervention. We intentionally design loops that have a human available to train our systems. These representatives increase our confidence that an AI solution is ready to go directly to the customer. We use this as an indicator to determine how quickly we are learning and training the system and to discover where it may need guardrails.

Projects that successfully move these leading indicators often go on to produce lagging indicator success as well. We use a three-stage funnel to move from estimation to validation to certified impact, comparing results with historical data at each step.

Support frameworks
With customer value and ROI as our north star, we found that we needed frameworks to align our teams, providing guidance for building in the short term while moving towards a long-term aspiration. These frameworks delineate what exists today

as well as what should exist in the future, mapping out a coordinated path to get there while keeping us true to our north star.

For example, the customer success flywheel clarified how deeply we needed to understand our customers in order to provide them with optimal value. The four-square framework maps customer value and human potential, offering a methodology by which to identify, develop, and prioritize opportunities.

Tips on leveraging AI
With the growing number of AI capabilities and business applications, people are often tempted to leap in to avoid missing out. This impulse leads many organizations to spend considerable time trying out various technologies and solutions. To reach the market briskly with maximum impact, we recommend three steps:

1. Identify the most important uses before picking the technology.
2. Treat AI solutions as building blocks that you can assemble to solve a bigger problem.
3. Develop prototypes with simple solutions with a focus on learning.

Our effort to manage the exponential demand on our marketing, sales, and support teams led us to AI-powered initiatives. We optimized each experience for different business goals, all developed with our customer-first principle which helped us to deliver the best to our customers. We used the prioritization framework as a second, important guide, challenging teams to push AI past the conventional boundaries of human ability and to prevent their pursuing easy wins that do not prioritize our customers' goals.

Our journey will never be done. We continue to look for new ways to deliver more value to our customers. ∎

Author Bios

The **Google Customer Engagement** team of authors consists of fourteen accomplished professionals. **Tim Frank** (lead author) has held various product management leadership positions in ads over the last ten+ years. His expertise includes computer/human interaction and human incentive systems, both of which are foundational to improving customer engagement. **Ravi Narasimhan, Qiushuang (Autumn) Zhang, Roman Karachinsky, Tony Li, Vicky Ge, Abheek Gupta, and Sandeep Beri** bring invaluable insights as leaders in product management. **Aastha Gaur and Doris Neubauer** provide a strong user experience perspective, aligning customer needs and business outcomes. Engineering leads **Leo Cheng** and **Ian Suttle** lend their expertise in machine learning, intelligent applications, and automated system architecture. **Natalie Mason** leads communications while **Shashi Upadhyay** is our general manager.

Endnotes

1. CSAT is measured via email and we interpret summed responses "happier than neutral" as satisfied.
2. https://www.deepmind.com/research
3. https://research.google/teams/brain/
4. For example, multi-task unified model (MUM) and language models for dialog applications (LaMDA).
5. Specifically 95 percent CI 3.9pp to 6.7pp
6. https://en.wikipedia.org/wiki/Long_short-term_memory
7. Specifically 95 percent CI 42.1 percent to 45.9 percent
8. Specifically 95 percent CI 46.1 percent to 49.9 percent

©shutterstock.com/Fit Ztudio

Singapore's Hospital to Home Program:
Raising Patient Engagement Through **AI**

John Abisheganaden
National Healthcare Group

Kheng Hock Lee
SingHealth Group

Lian Leng Low
SingHealth Group

Eugene Shum
SingHealth Group

Han Leong Goh
Integrated Health Information Systems

Christine Gia Lee Ang
Integrated Health Information Systems

Andy Wee An Ta
Integrated Health Information Systems

Steven M. Miller
Singapore Management University

Because of their complex care needs, many elderly patients are discharged from hospitals only to be readmitted for multiple stays within the following twelve months. John Abisheganaden and his fellow authors describe Singapore's Hospital to Home program, a community care initiative fueled by artificial intelligence.

Elderly patients, especially those who suffer from multiple chronic ailments and have complex care needs, often find it difficult to take care of themselves after they leave the public hospital where they have stayed for several days. As a result, many are readmitted over the next twelve to twenty-four months.[1,2]

To reduce readmissions, the Singapore public healthcare system devised the Hospital to Home (H2H) program, launched in April 2017, providing a nationwide community care framework that offers post-hospital help to this special group of patients, many of whom are frail and in need of additional medical and social support. The H2H program helps at-risk patients safely transition back to their home and community in a timely manner.

But how to predict who should be in the program? Many of the patients to be screened have complex post-discharge support needs and a high risk of returning to the hospital, not just once but several times over the subsequent twelve months. The Ministry of Health therefore deployed an artificial intelligence (AI) prediction model, trained on admissions and demographic data, to perform the initial screening of patients to determine whether they should be enrolled in the H2H program.

The AI prediction model trained on data from the National Electronic Health Record (NEHR) system, which flowed into the Business Research Analytics Insight Network (BRAIN) platform – a centralized business intelligence, analytics, and AI processing platform serving a wide range of data analytics needs across Singapore's public healthcare system. Within BRAIN, a scheduler system triggers the prediction model to run daily for all patients recently admitted to public hospitals, generating a patient readmission risk score for each.

The data used to train the AI prediction model includes more than 1,000 indicators that fall into three categories: sociodemographic characteristics, past hospital use, and past medical conditions. These indicators include the patient's age, number of non-elective inpatient admissions, total length of hospital stays in the past two years, total number of specialist outpatient visits, emergency department visits in the past year, and much more.

The AI prediction model

The machine learning method used for the model is a gradient boosting algorithm. We considered it essential that we be able to clearly explain the model, so we did not use deep learning methods. The gradient boosting algorithm allows us to readily show the relative influence of each variable on the results. As a result, clinician experts can check if these relative influences are in keeping with their clinical judgement.

The prediction model then identifies patients, admitted within the past twenty-four hours to a public hospital,

©shutterstock.com/nafterphoto
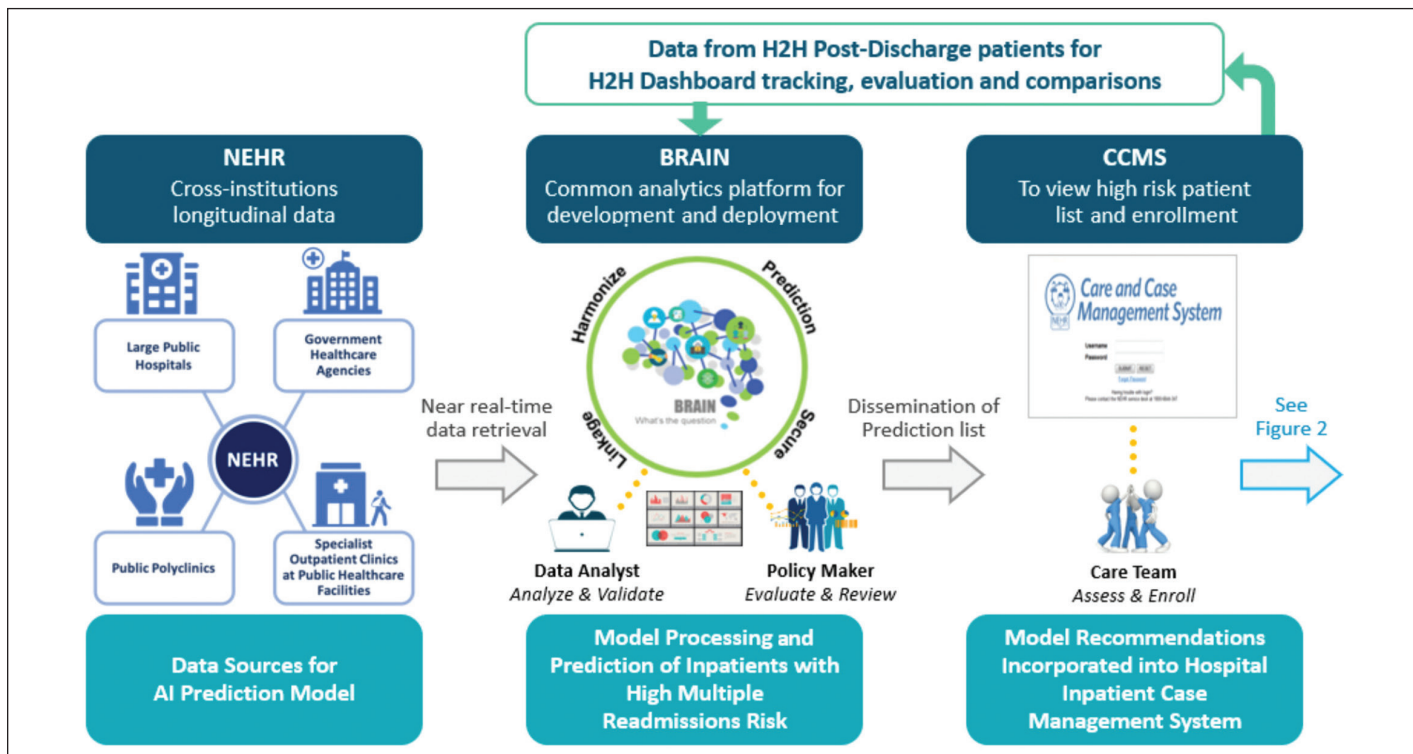
who are at high risk of multiple readmissions over the next twelve months. This result is incorporated into the hospital's care and case management system (CCMS) – a platform that provides a standardized view of each current patient and includes admission details, demographics, and past medical intervention history in the public health system.

After hospital staff use the predictive model to recommend patients for the H2H program, they vet them personally, using their clinical judgement to narrow the list of candidates for the program. Once chosen, patients who agree to enroll in the H2H program receive counseling and education about the program during the remainder of their hospital stay.

After they leave the hospital, an interdisciplinary team including doctors, nurses, allied health professionals, social workers, and community care coordinators provide these patients with further care and support through home visits and phone follow-ups.

Impact on patient engagement

Without the AI prediction model, the nurses who screen for for the H2H program in each public hospital would have to spend about half of every day manually screening the entire patient list and visiting wards to assess newly admitted patients.

This process would be impractical and too time intensive to carry out every day, in every ward of every public hospital. One large public hospital estimated that, without the AI prediction model, it would need fifteen nurses, rather than five, to complete daily screenings for the H2H program. The cost would be prohibitive.

The AI prediction model lightens that workload by identifying the 10 to 15 percent of patients who are at high risk of multiple readmissions over the next twelve months. The screening nurses, as well as doctors and other clinicians in the program, can then focus on vetting this smaller group without having to screen the other 85 to 90 percent.

Singapore's public hospitals admit more than 450,000 patients every year, around 40 percent of them aged 65 and older.[3] Using the AI prediction model reduces the average daily vetting workload from more than 1,200 total patients per day to an average of between 123 and 183 per day (10 to 15 percent of the national average of daily admissions), making screening much more manageable.

In addition, the AI prediction model's faster selection process gives the H2H multidisciplinary care team more time to engage with patients during their hospital stay. This additional engagement time is crucial for preparing patients and their families for what to do after leaving the hospital.[4]

Figure 2: Hospital to Home Patient Assessment, Care Plan Creation and Post-Hospital Support Process
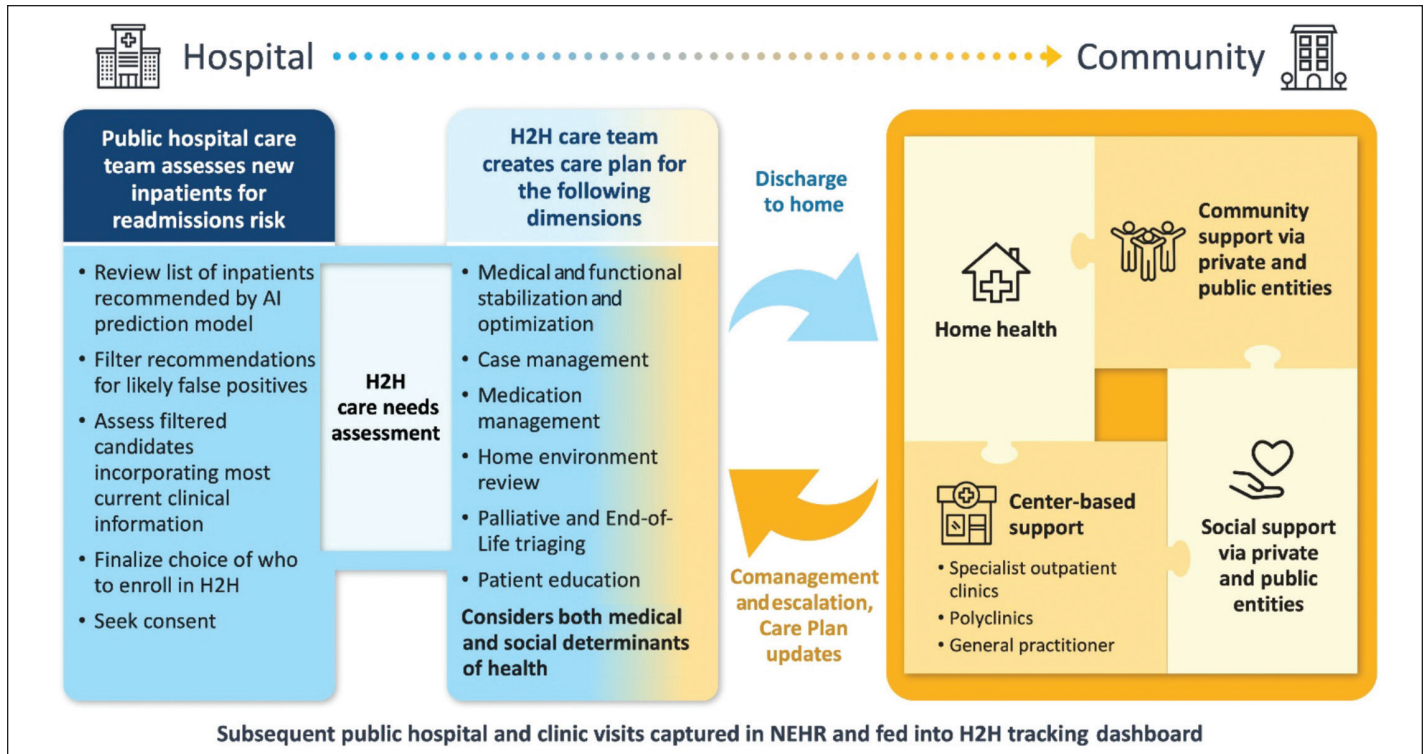
Without the AI prediction model, it would have taken three times longer for the current number of nurses to do the initial screening of patients for the H2H program. The second step of clinical vetting would then be delayed by two to three working days, which in turn would reduce the time left to explain to patients and their families what to do after discharge.

Doctors, nurses, and other members of the H2H care team appreciate the lighter workload and this savings in time, which lets them focus on the patients – both inpatients who are being enrolled into the program and discharged patients in the community care network.

Downstream benefits
The impact of the AI prediction model goes far beyond the H2H candidate selection process. Patients in the program tend to have reduced hospital stays in the 180 days after enrolling. Between April 2019 and April 2021, the average number of bed days saved across three of our major hospitals ranged from three to seven days per H2H patient.

To calculate the number of bed days saved, we compared patients who were enrolled in the H2H program to an unenrolled group who were of a similar age and Charlson Comorbidity Index (CCI) score.[5,6] We compared the total number of bed days in the 180 days after an index admission with those in the 180 days prior.[7]

> Without the AI prediction model, it would have taken three times longer for the current number of nurses to do the initial screening of patients for the H2H program

As clinicians, we know that shorter hospital stays benefit both the patient and the hospital.

Lessons for engagement
The AI prediction model allows human clinicians to efficiently and effectively engage with customers, in this case the hospital inpatients. After the enrolled patients are discharged and return home, they are assured close engagement with their H2H care team, which includes home visits from social workers and community health workers as well as phone and video conferencing with medical practitioners. This close interaction improves patients' health results and reduces hospital stays by avoiding future readmissions. This AI prediction model works behind the scenes to support human workers, and by doing so, allows our healthcare workers to engage directly and more deeply with their patients. ■

Author Bios

John Abisheganaden is Head of the Respiratory and Critical Care Department at Tan Tock Seng Hospital, National Healthcare Group, Singapore. He served as cochair of the national working group for the Hospital to Home program. He is part of the hospital's Complex Care Workgroup, focusing on coordinated care of patients with multiple specialist outpatient clinic appointments. He is an Adjunct Clinical Associate Professor at NTU's medical school and a Clinical Associate Professor at NUS's medical school. john_abisheganaden@ttsh.com.sg

Kheng Hock Lee is Deputy CEO (Education and Community Partnerships), SingHealth Community Hospitals, SingHealth Group, and Senior Consultant, Department of Family Medicine and Continuing Care, Singapore General Hospital. He is an Associate Professor at Duke-NUS Medical School. He set up the first hospital-based family medicine department in Singapore and led the implementation of the first hospital-based home medical service. lee.kheng.hock@singhealth.com.sg

Lian Leng Low is Director of the Population Health and Integrated Care Office, SingHealth Group, and Director of SingHealth's Centre for Population Health Research and Implementation. He is a Senior Consultant at the Department of Family Medicine and Continuing Care at Singapore General Hospital, Medical Director at Outram Community Hospital, and an Associate Professor at Duke-NUS Medical School. He is Director of Population Health Research for the SingHealth Duke-NUS Family Medicine Academic Clinical Program. low.lian.leng@singhealth.com.sg

Eugene Shum is Chief Community Development Officer, Changi General Hospital, and Director of Community Partnerships, SingHealth Group, Singapore. He led the development of Neighbors for Active Living, a precursor to the H2H program, that supported residents with high care needs in the community. He also oversaw the development of CareLine, a nationwide 24/7 urgent assistant telecare service for seniors. He is an Adjunct Associate Professor at Saw Swee Hock School of Public Health, National University of Singapore. eugene.shum.j.w@singhealth.com.sg

Han Leong Goh is Senior Principal Specialist, Data Analytics and AI Department, Integrated Health Information Systems, Singapore, heading the AI and Data Science team. He led the data science team modeling effort for developing the machine learning model used by the H2H program to predict multiple readmissions. He is currently involved in numerous other AI model development and deployment efforts. He is an Adjunct Senior Lecturer in the Analytics & Operations Department, NUS Business School. hanleong.goh@ihis.com.sg

Christine Gia Lee Ang is Director of the Services Planning, Data Analytics, and AI Department, Integrated Health Information Systems, Singapore. She leads IHiS efforts to plan and manage national analytics and AI projects, working closely with stakeholders from the Ministry of Health and public health institutions to explore, adopt, and exploit data and AI as part of their transformative programs to improve the quality of healthcare service delivery and population health. christine.ang@ihis.com.sg

Andy Wee An Ta is Chief Data Officer and Director of Data Analytics and AI for Integrated Health Information Systems, and Director, Health Insights, Singapore Ministry of Health. He is responsible for data analytics and AI-related services planning, data management, engineering and operations, model development and deployment, and all data governance for IHiS, and for envisioning, planning, and implementing data, analytical, and AI capabilities and programs across Singapore's public health ecosystem. andy.ta@ihis.com.sg

Steven M. Miller (corresponding author) is Professor Emeritus of Information Systems, School of Computing and Information Systems, Singapore Management University. Formerly, he was the school's founding dean, and the university's vice provost of research. He is coauthor of *Working With AI: Real Stories of Human-Machine Collaboration* (MIT Press, 2022) and in recent years has authored or coauthored numerous write-ups on AI applications and related digital transformation efforts across Singapore's public and private sectors. stevenmiller@smu.edu.sg

Endnotes

1. G. Facchinettia, D. D'Angelo, M. Piredda, T. Petitti, M. Matarese, A. Oliveti, M.G. De Marinis, "Continuity of care interventions for preventing hospital readmission of older people with chronic diseases: A meta-analysis," *International Journal of Nursing Studies*, January 2020, Vol 101, Article #103396. https://www.sciencedirect.com/science/article/abs/pii/S002074891930197X

2. Y.H. Ang, M.L. Ginting, C.H. Wong, C.W. Tew, C. Liu, N.R. Sivapragasam, D.B. Matchar, "From Hospital to Home: Impact of transitional care on cost, hospitalisation and mortality," *Annals, Singapore Academy of Medicine*, October 2019, 48(10):333-337. https://annals.edu.sg/pdf/48VolNo10Oct2019/V48N10p333.pdf

3. Ministry of Health, "Admissions and Outpatient Attendances." https://www.moh.gov.sg/resources-statistics/singapore-health-facts/admissions-and-outpatient-attendances

4. The importance of having as much time as possible within the current inpatient stay to educate and counsel the patient and their family circle about H2H post-discharge community care procedures is explained in Ta and Goh et. al., "Two Singapore public healthcare AI applications for national screening programs and other examples," *Health Care Science*, 19 August 2022.

5. M.E. Charlson, P. Pompei, K.L. Ales, C.R. MacKenzie. "A new method of classifying prognostic comorbidity in longitudinal studies: development and validation," J Chronic Dis. 1987;40(5):373-83. https://www.sciencedirect.com/science/article/abs/pii/0021968187901718?via%3Dihub

6. W.P. Glasheen, T. Cordier, R. Gumpina, G. Haugh, J. Davis, A. Renda. Charlson Comorbidity Index: ICD-9 Update and ICD- Translation. Am Health Drug Benefits. 2019 Jun-Jul;12(4):188-197. PMID: 31428236. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6684052/

7. The formula used is as follows: Net Change in Bed Days = Average Change in Bed Days of Enrolled Patients – Average Change in Bed Days of Non-Enrolled Patients of the same CCI score (control group); where Change in Bed Days = Cumulative Bed Days 180 days post the Index Admission– Cumulative Bed Days 180 days prior to the index admission.

Commonwealth Bank: Amplifying Customer Centricity with AI

Steven Randazzo, Warwick Business School and Harvard University
Jin H. Paik, Harvard Business School and Altruistic
Yael Grushka-Cockayne, Darden School of Business, University of Virginia

Steven Randazzo, Jin H. Paik, and Yael Grushka-Cockayne describe how Commonwealth Bank used data and AI to maintain a competitive edge by bridging gaps between retail, call center, and digital services. By making data more available and standardized, the bank improved customization and enhanced its interactions with customers.

Success in retail banking hinges on retaining customers in order to capture their lifetime value. Until now, many banks have succeeded in attracting new customers through a one-size-fits-all marketing approach. The efficacy of this traditional approach has led retail banks to also apply it to retaining customers, investing marketing dollars in both traditional and digital channels.

Like other banks, Commonwealth Bank, or CommBank, Australia's largest retail bank, with 15 million customers, took this approach. "It was a fairly heavy sales-oriented culture, with a largely homogenous approach to customer interactions," explained Matt Malady, general manager for channel and customer at CommBank.

But with a changing customer landscape, CommBank decided to rethink its approach rather than simply funneling more resources into the existing strategy, especially since it had access to large amounts of data. "In the future, our customers will increasingly be able to serve themselves rather than coming to us, so we have got to find a way to be more relevant and personalized in meeting their needs than we ever have been before," said Angus Sullivan, group executive of retail banking.

The leaders of CommBank decided to shift to an artificial intelligence-powered approach to creating new value. This approach would use data to augment the abilities of employees, transforming every frontline worker into a customer expert, saying the right thing, at the right time, to every customer.

The CommBank Approach

To fundamentally shift its approach to customer engagement, CommBank tapped into the data and business expertise of the organization, using existing components to create a new framework.

But the new approach did not match the existing technical and operating structures of CommBank. Its three primary channels – branch, call center, and digital – were poorly coordinated, which led to duplication and confused messaging and supported the one-size-fits-all approach.

Yet these three channels included more than 1,000 branches and 48,000 employees, 2,000 of them in customer contact centers, making and receiving calls. CommBank's digital footprint consisted of its website and popular mobile app.

AI-based digital transformation meant changes to infrastructure, workforce, and organizational structure.

AI-based digital transformation meant changes to infrastructure, workforce, and organizational structure. And cooking up the secret sauce of creating value was difficult. "How do we, over thousands of interactions, try to generate the same outcomes as from an in-depth, one-to-one conversation?" Sullivan asked.

To develop this new approach, CommBank tapped specialists who had knowledge of the different pieces that make up an AI system, as well as leaders and experts with architectural knowledge of the existing framework and what needed to be changed. This team was charged with linking all the components together into a new, coherent whole.

The result was the Customer Relationship Banking Program (CRB), an internal program intended to redesign the customer experience throughout the bank. Its flagship project was a customer engagement engine (CEE), an AI-driven customer-experience platform.

The Sweeping Impact of CEE

It is no hyperbole to say that the CEE transformed how CommBank employees interacted with customers. The CEE makes 35 million decisions per day, devoting less than 300 milliseconds to each. It provides full cross-channel integration using 450 machine learning models that learn from a total of 157 billion data points.

For example, when customers go into a branch, the CEE suggests options for a next best conversation (NBC) to the workers who greet them, ranking and scoring possible messages. The staff can then use these conversation starters to improve the customers' experience.

The system first assesses whether a customer should receive a specific message. Then it evaluates the NBCs, determining which are appropriate to the context. Finally it presents its selections to the frontline staff as options. Employees then use their judgment to select an NBC. After the interaction, employees record feedback on how the customer responded to the NBC. The bank has applied this process across all channels, and customer behavior determines which NBCs are most relevant. (See Figure 1)

The success of an NBC is measured by the NBC acceptance rate, that is, what the customer does after being presented with an NBC. For in-person interactions, tellers note in the system how the NBC was received; online, the NBC acceptance rate is measured by what the user does next. CommBank uses this acceptance rate to modify its NBCs to better fit customers' needs.

In one example of the value an NBC can have, a teller was prompted to wish a customer a happy 90th birthday. The customer was delighted. "How did you know it

was my birthday? Since I live alone, you are the first person to wish me a happy birthday."

The CEE and Customer Engagement

The CEE eliminates the ambiguity of manual assessments as well as the expertise necessary to overcome it. Applying principles of behavioral economics, the bank sent various NBCs to customers to test whether the CEE could help them make better decisions about using their tax refunds than previous campaigns. It achieved a fourfold increase in NBC acceptance over years past.

In another experiment, Comm Bank offered a 'fuel finder' to 250,000 New South Wales customers, providing current information on the cheapest fuel in their area as well as a price map. Customers could choose to receive NBC alerts to help them save on their preferred fuel type. The system sent messages by using past fuel spending patterns to predict when the specific customer would need to refuel their vehicle. Over two test months, the bank's acceptance rate was ten times that of the standard message suite.

For the contact center, the CEE identifies customers who regularly call and sends them messages offering more convenient options. While on the phone, staff members receive a push notification helping them to direct the customer into chat messaging. More than 40 percent of call center interactions are now performed via messaging, including those performed in the bank's app.

The CEE allows CommBank to make best use of its data, identify customers instantly, and offer immediate support that matches their needs. The CEE's groundbreaking ability to improve the performance of frontline staff, making everyone an expert to improve the experience of customers, puts it far ahead of the systems of comparable banks.

Figure 1. Sample Next Best Conversations (NBCs).

The CEE also raised the bank's annual net promoter score (NPS). In the 2020-2021 period, the home lending teams saw a tenfold increase in leads that were 300 percent better than those of previous years. This improvement resulted in more conversions and an increase in the NPS score among mortgage customers by 16.4 points, to a score of 10.4. Improvements in NPS also spilled over to CommBank's digital channels, with 1.2 million more customers logging in daily and more frequently, up from 32 to 34 times on average per month.

In 2019, the CEE delivered its 50 millionth NBC to a frontline employee, transforming their interaction with a customer. The bank's leaders have determined three factors, which other companies can readily adopt, to be crucial to success.

1. **Explainability:** CEE managers include a message on the NBC screen that explains why the suggested message is relevant to the customer and situation. Organizations can replicate CommBank's approach by making it clear why users are receiving a particular prediction, which reduces the uncertainty often associated with AI 'magic.'

2. **Training:** CEE managers developed a training program for staff to prepare them for the drastic change in operations. Managers made more experienced employees advocates of the process, asking them to coach other team members. The goal was to integrate the messaging into everyday speech without sounding scripted.

CEE managers also set up a feedback forum to give employees confidence in their modified duties while also allowing them to share their feedback about the approach and the NBCs. Managers used this feedback to improve NBC development, honing the suggested messages. Organizations should take staff feedback into account to ensure that their algorithms make correct predictions and to establish new pathways for data collection to further inform the AI.

3. **Content:** The CEE team carefully considered which messages would resonate with customers. Should they start with sales messages that could generate revenue or with informative messages? The pilot program included some sales messages, but ultimately the team decided to favor service messages, such as notifying people of their anniversary with the bank, thanking them for their fidelity to the bank, and so forth.

Customers responded well to this approach, which helped frontline staff to feel confident in the technology. As one CommBank manager put it, "It is an active way to show that we care."

Since that 50 millionth message, CommBank's use of CEE with customers has continued to evolve, moving beyond service and sales messages to new programs that help and support customers. The CEE can now connect customers with financial relief programs during difficult times like the COVID-19 pandemic, wildfires, and other disasters, as well as during economic shocks like persistent inflation.

> The bank moved from a siloed, generalized approach to one that was data-driven and personalized.

CommBank continues to use the CEE as a core piece of its customer experience strategy. With the development of the CEE, the bank moved from a siloed, generalized approach to one that was data-driven and personalized, translating millions of customer data points and combining them with the expert knowledge of the frontline staff.

The success of this transformation was not a given; it hinged on key investments in infrastructure, talent, and a culture that embraced learning and AI. Human plus machine created an improved experience that focused on creating value for the customer by assessing, predicting, and responding to their needs while capturing value for the bank. ■

Author Bios

Steven Randazzo is a PhD student at Warwick Business School and a visiting research fellow at the Laboratory for Innovation Science at Harvard. Steven's research focus is on the use of AI in firms and how AI is creating new value and capturing opportunities in complex knowledge work. Steven is also Head of Innovation for Altruistic, an AI and data science consultancy.

Jin H. Paik is a research scientist at the Harvard Business School and a cofounder and managing partner of Altruistic, a data science consultancy. He was the head of labs at the Data, Digital, and Design Institute at Harvard, where he developed the institute's strategic vision and directed project and research activities. He earned a bachelor's degree from the University of Michigan and a master's degree from Harvard and is a doctoral candidate at New York University.

Yael Grushka-Cockayne is a professor of Business Administration, the Altec Styslinger Foundation Bicentennial Chair in Business Administration and Senior Associate Dean for Professional Degree Programs at the University of Virginia, Darden School of Business. Her research and teaching focus on data science, forecasting, project management and behavioral decision making. She is published in numerous academic and professional journals and is a regular speaker at international conferences on decision analysis, project management, and management science.
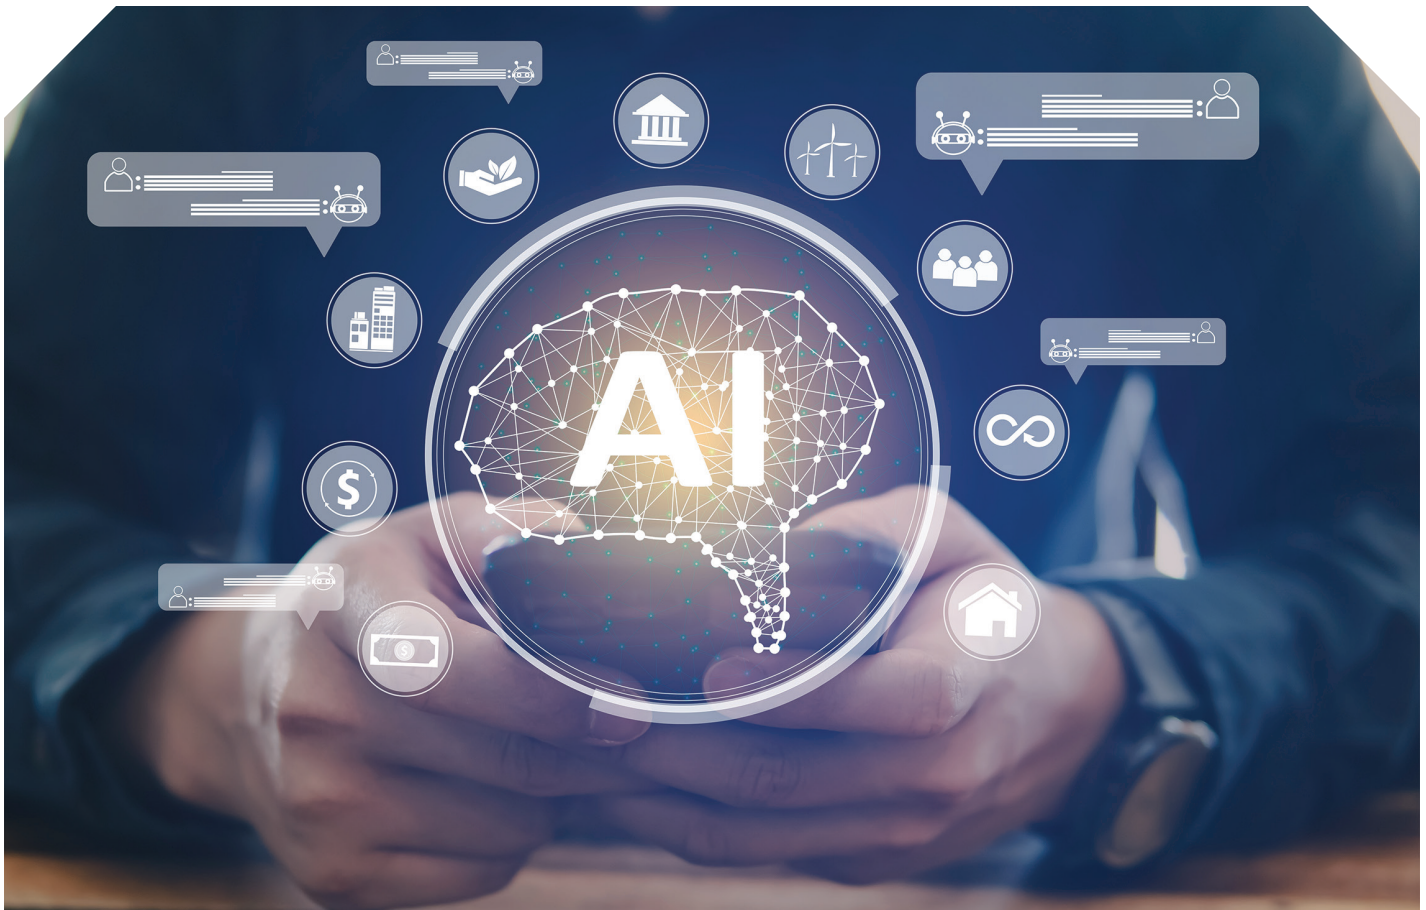
Singapore's AI Applications in the Public Sector: Six Examples

Steven M. Miller, Singapore Management University

Steven M. Miller describes six instances in which Singapore has applied AI in the public sector, illustrating different ways of improving its engagement with the public by making government services more accessible, anywhere, anytime, and speeding its responses to public processes and feedback. He illustrates how its leaders made the city a living lab for AI use, and what they learned.

In November 2019, Singapore's Smart National and Digital Government Office (SNDGO) unveiled its National AI Strategy, stating its ambition that, by 2030, the country would be a leader in developing and deploying scalable, impactful artificial intelligence (AI) solutions in key economic and social sectors of high value and relevance, and that it would do so in ways that would serve the needs of its citizens and businesses.[1] These efforts would be concentrated in seven domains deemed crucial for a smart nation.[2]

Six of Singapore's applications of AI in the public sector illustrate different ways that AI an be used to enhance engagement with the public. Table 1 summarizes these efforts.[3]

The Six Examples

1. Managing Crowd Levels at Parks

Singapore's National Parks Board, known as NParks, oversees the city's various parks, gardens, and nature reserves, striving to make Singapore "a city in nature."[4] Providing and maintaining park access is an important public service that is especially appreciated in a small city-state with a high population density.

> NParks needed to not only help its staff to manage this situation, but also encourage members of the public to make better and more informed choices about which park to visit and when.

In early 2020, COVID-19 restrictions drove more people to outdoor park areas. This sudden surge required that NParks staff be diverted from their regular work to count visitors and to patrol and enforce safe distance regulations in the more than 500 green spaces they managed. NParks needed to quickly come up with a solution that would not only help its staff to manage this situation, but also encourage members of the public to make better and more informed choices about which park to visit and when.

NParks worked with the Government Technology Agency of Singapore (GovTech) to create an AI-enabled system called Safe Distance @ Parks. The tool combines input from the parks' video streams with data about car park occupancy to determine how many people are in green spaces across Singapore in near real time (Figure 2).[5] However, not all of the video cameras already installed at parks were equipped to livestream to a central cloud platform for analysis. The video coverage was also insufficient to meet the operational needs of this new system. In addressing these deficiencies, NParks and GovTech decided to prioritize the more popular parks.

In just two weeks, GovTech deployed five video cameras, building on its ongoing projects and existing platforms. The team could then rapidly verify its methods for transmitting all images and other data back to a central cloud platform for analysis, its choice of models to achieve the necessary accuracy for counting people, and its methods for updating its numbers.

The rapid prototype also allowed the team to determine all necessary protocols and APIs (application programming interfaces), and to integrate the system with the Safe Distance @ Parks website and mobile app that would be used by NParks staff and the public.[6] The first version of the public website was launched a week after the technical demonstration, and then progressively refined by the team. Once all the technical issues were worked out, GovTech and NParks spent the next eight months refining the system and expanding its coverage to more park sites.

They also experimented with drones and mobile robots to safely give officers a better sense of visitor density in specific areas, allowing for more targeted safe distancing efforts. In a few park locations, especially during periods of heightened COVID-19 concern, the team used Spot, the robotic dog, to broadcast a recorded message

Figure 1: Seven domains for National AI projects
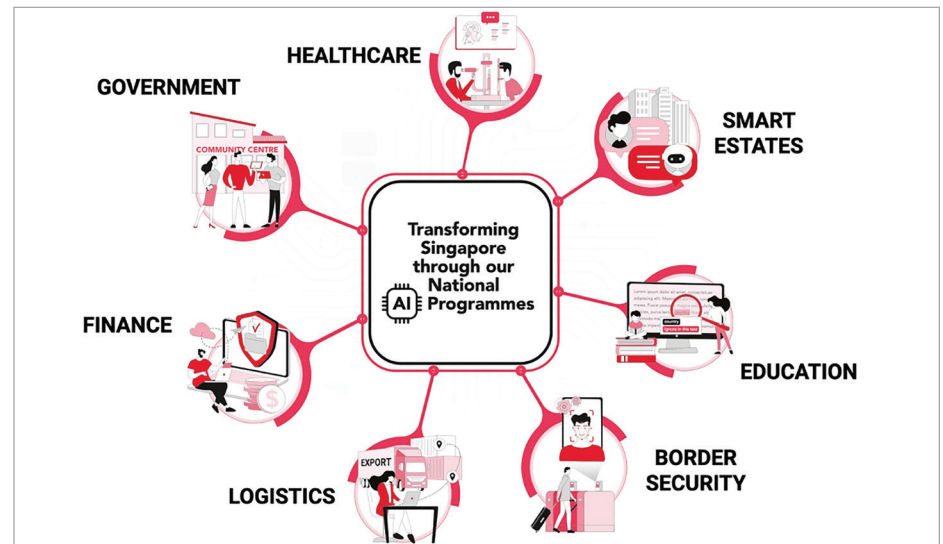
Source: Smart Nation and Digital Government Office (Singapore)

Table 1. Summary of Singapore's public sector AI applications			
	Public Sector	**AI Application**	**Impact on Public Engagement**
1	Government, Everyday Services	Managing the number of people in parks during COVID-19	▪ Eliminates tedious crowd-counting by National Park staff, frees them to do other park work and public engagement. ▪ Web portal lets the public check park crowd conditions.
2*		Job search and matching in support of future transformation of the economy	▪ Reduces job mismatches and missed matches for both local job seekers and potential employers, uses contextualized local labor market information and data. ▪ Expands the range of recommendations to include adjacent job options and educational courses to introduce new career possibilities.
3		Mass transit, urban rail system management with a customer-centric monitoring approach	▪ Faster and better response to service disruptions and crowded commuter conditions across the entire urban rail network.
4*	Smart estates, seamless and efficient municipal services	Municipal services office and the OneService chatbot	▪ Easier to submit complaint reports on municipal issues via a new chatbot system with a more flexible interface and more intelligent backend processing capabilities.
5*	Education, K–12	Adaptive learning and assessment for personalized public schools	▪ AI systems that personalize the education of each student in selected subject areas. ▪ Teachers can shift towards more customized and targeted engagement with both individual students and groups.
6*	Border Security	Contactless (breeze-through) immigration clearance at border control checkpoints	▪ Shorter wait times for travelers at immigration checkpoints; less physical contact with documents or equipment for travelers and checkpoint staff. ▪ More consistent, evidence-based approach to identifying travelers in need of special screening.

*Applications designated as National AI projects falling under the National AI domain areas in Figure 1.

Figure 2: Safe Distance @ Parks solution architecture overview
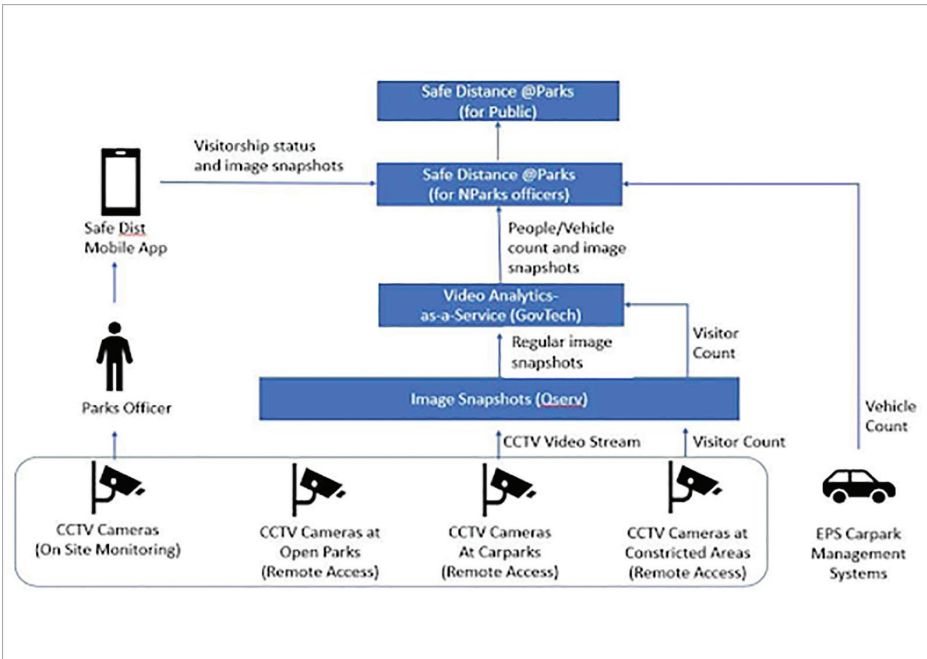
Source: GovTech (Singapore)

reminding visitors to observe safe distancing measures.[7,8]

Spot was also outfitted with cameras with GovTech video analytics to estimate the number of visitors it encountered (Figure 3). To allay privacy concerns and build public acceptance, the cameras and image analysis did not track or recognize specific individuals and no personal data was collected. NParks and GovTech continue to refine the behavior of the robotic dog to bring greater nuance to its interactions with humans in the park and to increase public acceptance of this approach.[9]

NParks made the Safe Distance @ Parks application accessible to the public and created a convenient way for anyone heading out to a park to see the crowd levels at that particular green space in advance.[10] People could then decide

where and when to go. Meanwhile the team continued to improve the public website to include real-time crowd heat maps of each park, seasonal crowd data, and other public leisure and recreational spaces.

At the height of COVID restrictions in 2020, monthly traffic to the Safe Distance @ Parks app peaked at 620,000 visits per month. Towards the end of 2021, this traffic had stabilized to around 400,000 visits per month. The app had automated people-counting, reducing the need for human enforcement staff by two thirds.

2. Local Job Search and Matching

Even before COVID-19, the Singapore government encouraged and supported working adults to upskill and reskill in response to technology changes, digital transformation, and ongoing industry changes.[11] All segments of the population were encouraged to adopt the habits and mindset of lifelong learners.[12]

The government, through Workforce Singapore (the former Workforce Development Agency under the Ministry of Manpower) launched the MyCareersFuture job portal to help residents find jobs, review career options and explore new job opportunities.[13] The self-help online portal also matches local job seekers with local employers.[14]

Workforce Singapore launched the first version of the jobs portal in April 2018, using machine learning and other AI methods to find the best job matches for each jobseeker's skills. Each match is quantified with a score from zero to one hundred; the higher the number, the closer the match. This score is drawn from an analysis of job descriptions and necessary skills, relative to the profile of the job seeker. The system also filtered results to include jobs with government training support schemes.[15]

Figure 3: Field trial of Spot, the robotic dog, supporting enforcement of social distancing in a Singapore park during the early days of Covid-19

Source: Photo by Roslan Rahman/AFP via Getty Images from a Los Angeles Times story

A year later, the team enhanced the MyCareersFuture system to make it easier for employers to sort and list applicants, and to preview job seekers.[16] Machine learning enhancements allowed the system to recommend potential job candidates, including those who met the requirements but applied to jobs in other areas.

With the economic upheaval caused by COVID-19, the active job postings on the MyCareersFuture portal (both employers looking for workers and people looking for jobs) nearly doubled, from about 87,000 in November 2020 to about 150,000 in October 2021. This change reflected personnel shortages for both online services and frontline COVID support, as well as government train-and-place schemes intended to help workers move into the technology sector. In the ensuing months, job listings subsided to about 130,000. During the same period, the portal

Figure 4: Overview of the JumpStart platform used by MyCareersFuture and other government apps requiring job matching recommendations
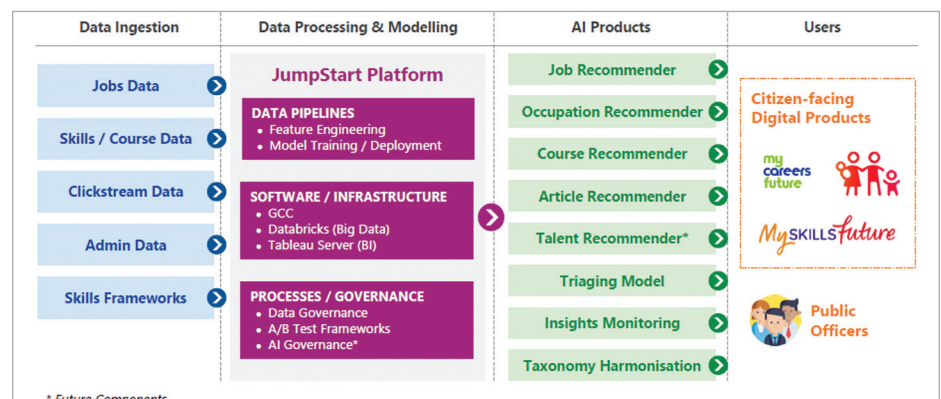
Source: GovTech (Singapore)

received more than 560,000 job applications, making it clear that Workforce Singapore needed to improve the portal's capabilities. By creating a new AI job-matching platform called JumpStart (Figure 4),[17] the team improved MyCareersFuture in four keyways:

A. A more flexible, microservices-based architecture

JumpStart was designed as an independent, cloud-based, centralized job search and matching platform linked to MyCareersFuture. This architectural approach made it easy to reuse Jumpstart for other government applications that required a job search and matching engine. Applying the microservices architecture to the job portal also made it easier to integrate functions from commercial AI vendors into JumpStart's capabilities.

B. Better methods for filtering, matching, and recommending

An earlier version of MyCareersFuture deluged job seekers with as many as 1,000 openings while directing a torrent of applications at employers. JumpStart uses a more finely tuned AI that matches workers and employers more closely. The improved design also incorporates job openings that are highly or overly subscribed, more closely identifies adjacent occupations to which a worker did not apply, and prioritizes local jobs and industries with good growth prospects. JumpStart also suggests continuing education courses.

C. Seeing beyond the job seeker's stated preferences to implicit signals

Beyond using explicit data from the resumes and profiles of job seekers, JumpStart also analyzes implicit signals, including past job searches, job views, job applications, which job recommendations seekers viewed, and for which job opportunities they applied.

JumpStart uses this data to develop several models, including:

 i. Skills-matching Model: A content-based filtering model that looks at how closely a job-seeker's skills match employers' requirements
 ii. Views-based Model: A collaborative filtering model that uses singular value decomposition to evaluate a jobseeker's previous views
 iii. Application-based Model: A collaborative filtering model built on jobseekers' previous applications that identifies potential jobs through similarity with other users

Because JumpStart's machine learning algorithms are trained on local data from the MyCareersFuture portal, its job search terminology, classifications, and recommendations reflect local labor market preferences, conditions, and trends. Singapore residents can also filter for listed job openings that are eligible for Workforce Singapore government support programs and easily link to the SkillsFuture listings of local training courses with government tuition subsidies.

D. Better support for using multiple data science models in parallel

JumpStart's job recommendations are generated using a hybrid approach in which multiple models (including the three mentioned above) are built and then mixed to work with one another to generate the eventual recommendations for the user.[18] The GovTech team did not want to build bespoke deployment solutions for each data science model. By integrating with the MyCareersFuture portal through a microservice architecture, the JumpStart team could experiment extensively with different ways of setting up a hybrid model and evaluating its per-formance without having to build and integrate bespoke deployments for each successive model change.

After GovTech deployed Jump-Start, the clickthrough rate at the MyCareersFuture portal doubled from 4 percent to 5 percent to 10 percent to 12 percent. Seven months after JumpStart's launch, more than 5,000 job seekers had secured jobs through Jump-start-enabled recommendations.

3. Mass Transit Urban Rail System Management

Singapore's state-of-the-art urban mass transit system makes extensive use of sensors to collect information, monitor operations, and gather data for AI models that predict disturbances. While the old system kept track of the current and predicted state of the rail network, it did not evaluate the experience of commuters. Singapore's Land Transit Authority (LTA) set out to monitor and improve the customer experience, using IoT sensing, situation assessment, and incident response planning.[19]

To monitor and manage these commuter-focused measures, which tracked how many passengers could not board because a train was too packed, the size of crowds on platforms and in trains, and the duration of delays, LTA and GovTech developed the Fusion AnalyticS for public Transport Event Response (FASTER) system, deploying it in mid-2018.

FASTER is an AI-powered data fusion platform that mines data from IoT sources, including video streams, WiFi and cellular signals, farecard data, train engineering and flow data, and taxi and other transport data to provide immediate warning of potential rail anomalies (Figure 5). The system makes the entire urban rail network visible around the clock, detecting unusual events and providing automatic alerts when it predicts impending disturbances.

Once a disturbance is predicted or occurring, FASTER gives rail operations staff enhanced real-time visibility at the micro and macro levels, allowing them to see how the event is affecting other rail stations so they can assess the impact system-wide. This visibility, combined with supporting real-time simulation abilities, lets the staff make thoroughly informed decisions.

FASTER also helps staff to quickly assess how various options would play out, whether adding trains to handle passenger backlogs, or, in more severe cases, deploying special buses to transport stranded passengers to unaffected rail stations.

FASTER has made the rail operations monitoring and response team more agile, productive, and efficient.

FASTER's outcomes are also fed into a broader LTA command, control, and communication (C3) system that draws status information and predictive insights about other land transport domains, such as roads and buses, to create a comprehensive view of the land transport network and coordinate fast and effective responses to large-scale disturbances.

When FASTER was first deployed, it was able to predict about 40 percent of impending incidents. By 2019, it had accumulated enough operating experience, fine tuning, and enhancement to predict 80 percent, and by the end of 2021 it approached 90 percent. This predictive ability allows rail operations staff to respond to most flow

Figure 5: Overview of the FASTER system
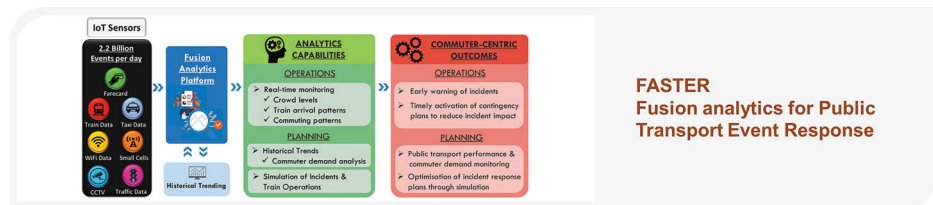
Source: Land Transport Authority of Singapore

Figure 6: Overview of the OneService chatbot
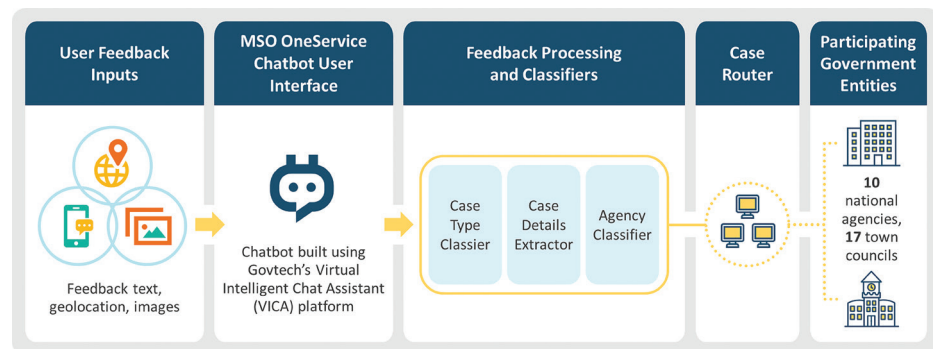
Source: GovTech (Singapore)

disturbances preemptively, usually minutes or more in advance.

The FASTER system has also increased the productivity of the rail operations monitoring team. At the end of 2021, a team of four monitoring officers per shift were managing twice the number of stations monitored by their 2012 counterparts. FASTER has made the rail operations monitoring and response team agile, productive, and efficient.

4. Municipal Services and the OneService Chatbot

In 2014, an infamous incident of a fishball stick discarded on a walkway highlighted the difficulties the public faced in reporting litter to the right public agency. In this case, three agencies fell to quibbling about their overlapping jurisdiction, none willing to remove the offending stick. That summer, the Prime Minister referred to the fishball stick incident in his National Day Speech as the driving force behind creating a new public agency called the Municipal Services Office (MSO). The MSO focused on

improving the government's overall coordination and delivery of municipal and neighborhood services.[20]

MSO's hotline call center soon became the one-stop-shop for citizens to report municipal and neighborhood problems. The MSO would then figure out which government agency was responsible and coordinate various agencies when necessary. Ten public agencies and seventeen neighborhood councils now work closely with the MSO.

MSO's OneService mobile app was launched in 2015, a community platform for all things neighborhood, where residents can provide feedback on municipal issues, learn more about their neighborhood, and access government and town council e-services. The call center was eventually discontinued, leaving the app as the only form of communication until the MSO developed a chatbot.

While the OneService app is widely used, in certain ways, it did restrict public engagement with the MSO.[21] Some people did not want to

download and use the OneService app, preferring messaging apps like WhatsApp or Telegram. And submitting a case report could be complicated for users, especially for those who did not understand the app's pre-defined categories, whose complaint did not cleanly fit into those categories, or who had more than one issue.

Nonetheless, the app collected hundreds of thousands of case reports that formed a large, diverse dataset of complaints and the agencies that resolved them. This collection made it possible to use data-driven AI approaches, rooted in machine learning.

In July 2021, MSO unveiled its OneService chatbot, named Kaki (the local Singlish and Malay expression for 'buddy' or 'companion') (Figure 6).[22] Through live chat sessions, Kaki can identify the nature of a complaint and place it in the appropriate category (e.g., litter, illegal parking, broken toilet, dead animal, etc.). It can also extract the relevant details of the incident and automatically generate the case report.

After the chat session ends and there is more time for computational processing, the backend system predicts the right government agency and automatically routes the report there. The system uses all the text from the chat as well as any photos or videos uploaded to make this decision.

Kaki's initial accuracy rates were 80 percent for categorizing the problem and 85 percent each for generating the report with all required attributes and assigning

Nine months after deployment, Kaki was handling over 300 user sessions per day, serving over 4,000 people.

it to the appropriate government agency. Accuracy rates stayed about the same in the following year, but Kaki expanded to handle many more types of complaints and many more users. Later, a more intelligent chatbot interface resulted in more than 90 percent of the cases being automatically routed to the right agency.

Nine months after deployment, Kaki was handling over 300 user sessions per day, serving over 4,000 people. Although these usage numbers were comparatively small, MSO and GovTech viewed them as evidence that the platform was effective. It was also seen as a successful application of GovTech's virtual intelligent chat assistant (VICA) platform, that standardized various government chatbots.[23,24]

5. Adaptive and Personalized Learning for Public Schools

One of the projects highlighted by Singapore's National AI Strategy was the use of AI to provide personalized education through adaptive learning and assessment for primary and secondary level public school students. The goal is to allow all students to learn at their own pace, letting them better absorb lessons while supporting teachers in meeting the personal educational needs of each student.

Here, GovTech used two types of AI-enabled solutions. An adaptive learning system teaches mathematics, using machine learning to offer advice on solving problems and responding to errors that is customized for each student. By analyzing the student's responses to the system's content and questions, it infers how well the student understands a concept. The engine then makes recommendations about what content and assessment items the student should move on to next.

Meanwhile, a learning feedback assistant for English language

writing gives students feedback on their grammar, sentence structure, and language errors. These two AI systems are delivered through the Ministry of Education's (MOE) Singapore Student Learning Space, the national online learning platform for grades 1 through 12.[25]

MOE piloted the AI adaptive learning system for primary and lower secondary school mathematics in cooperation with a few schools in 2022.[26] Two years earlier, MOE had shared the experiences of two secondary schools with a small-scale trial of commercial adaptive learning systems for some of their math classes.[27] Some teachers noted that they were able to use student performance and information from the adaptive learning system to encourage and facilitate peer teaching and learning.

Teachers could also more readily identify students who were struggling, offering them more attention and guidance. One teacher noted that some students responded to the different learning stages as if they were levels of a game and strove to keep climbing upward to the next level. These students appeared to enjoy the game approach, regardless of their math skill.

In August 2021, MOE called for bids for the development, piloting, and implementation of an AI system to provide feedback on English language writing. The system, to be deployed in 2023, will provide personalized feedback to improve students' English language writing skills by correcting their grammar, sentence structure, and expression.[28] An MOE official noted that it would later consider rolling out similar systems for other languages.

In November 2021, parliament questioned MOE about whether the use of this type of automated AI could make students' approach writing routine, undermining their ability to write and think with

creativity, cadence, personality, and flair.[29]

MOE responded that its use of educational technology, including this application, does not dilute, but complements the work of the teacher in guiding students to learn and master subjects, and especially high-level skills. By relieving teachers of the task of checking and correcting basic elements of writing, the system frees them to devote time with students to more complex aspects of language and advanced writing skills like creative expression, persuasiveness, and tone.

> AI support for border security and crossing management is one of the designated national AI projects.

6. Contactless Clearance at Border Control Checkpoints

Singapore is deliberate about not only maintaining but always improving its image as a welcoming and safe destination for both tourists and businesses. The challenge is to insure that all inbound and outbound travelers can clear immigration quickly and without hassle without compromising border security.

AI support for border security and crossing management is thus one of the designated national AI projects; the 2019 National AI Strategy states, "We aim to deploy AI to achieve 100 percent automated immigration clearance for all travelers, including first-time social visitors. Singaporeans and departing visitors will experience 'breeze-through' immigration clearance, without the need to present their passports."

Contactless immigration clearance, using AI to analyze iris and facial scans to verify a traveler's identity, is currently in use at Singa-

pore's land border checkpoints and at the airport (Figure 7). With a verified identity, backend systems can also verify passport status and other information relevant to crossing the border. Many countries perform facial matching using biometrics in the chip in the traveler's passport. Singapore chooses instead to use its own databases for biometric matching because it is more secure and not susceptible to forged passport chips.

In order to function, this system must capture the necessary iris and facial scans. The Immigration & Checkpoints Authority (ICA), an agency under the Ministry of Home Affairs, started encoding the iris and facial images of Singapore citizens and permanent residents in January of 2017, as part of the process of issuing, renewing, or updating passports or national identity cards.

Singapore citizens, permanent residents, and long-term pass holders, as well as international travelers in ICA's frequent traveler program who have

enrolled their iris and facial biometrics are eligible for the new breeze-through clearance.[30] First-time foreign visitors to Singapore can also use the breeze-through option if they enroll their entire set of biometrics at the border. If, for any reason, the system can't get an iris and facial scan, it prompts travelers to scan their fingerprint as a backup.

Breeze-through clearance makes it possible for ICA to handle the growing volume of travelers crossing Singapore's borders. After the rollback of many COVID-19 travel restrictions, border crossing volumes rapidly recovered. ICA kept pace by using AI-enabled technology to reduce the amount of time immigration officers spend on routine low-value tasks. These officers can instead focus on high value tasks such as interpreting analytics and following up with special interest travelers identified by ICA databases from international or domestic watchlists.

Home Team Science and Technology Agency (HTX), which is dedicated to supporting the needs of

Figure 7: Contactless breeze-through immigration clearance
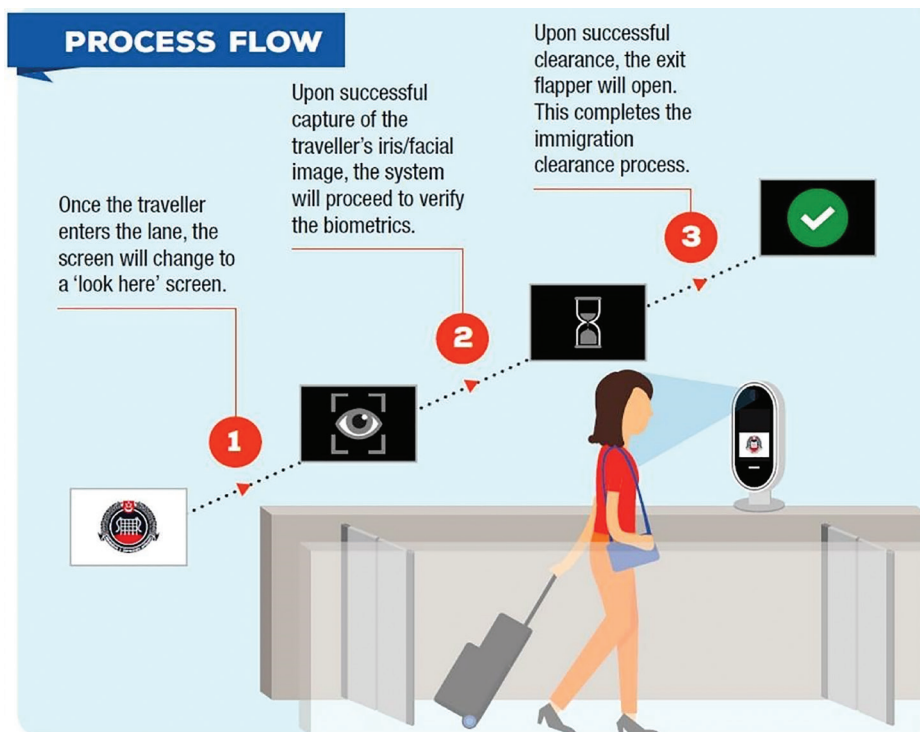
Source: Immigration & Checkpoints Authority of Singapore

the Ministry of Home Affairs, used a combination of purchasing and internal development for iris and face detection systems. Within HTX, staff from the Immigration & Checkpoints Programme Management Center and the Biometrics and Profiling Center of Expertise worked together to create, deploy, and continually improve contactless border checkpoint solutions and their supporting technology.[31]

HTX's Biometrics and Profiling (B&P) Center of Expertise hopes to improve the system's ability to detect suspicious travelers through the automated detection of the behavior and physiological reactions characteristic of people with malicious intent. The system would then translate tell-tale indicators (TTI) into prediction algorithms to identify and point out such behavior.

TTIs are usually associated with a person's observable behavioral and non-verbal cues, such as fidgeting or excessive sweating. By developing an automated TTI detector, ICA expects to increase the consistency and objectivity of profiling at border checkpoints and of decisions about which travelers should be more intensely questioned.

Already travelers using breeze-through have a contactless way to provide their biometrics for verifying identity. Eliminating the need to check fingerprints or handle passports speeds up border clearance. Faster processing time reduces the time that travelers spend waiting in queues at the border checkpoint. And immigration officers using AI screening and profiling receive guidance on the very small subset of inbound and outbound travelers who are candidates for additional security measures.

Conclusions

These six examples of the Singapore government's digital and AI journey have important lessons to teach:[32]

A. Three important dimensions of public engagement are accessibility, responsiveness, and personalization.

AI can improve public sector engagement with all three. While regular digital technologies can support 24/7 access, applying AI to these channels creates more intelligent retrieval, processing, and evaluation of large data from various sources, and also supports or automates the triaging of complex problems. This capacity results in more responsive and personalized interactions with customers.

AI fueled personalization also improves public access to relevant information so people can make their own informed choices as they do with Safe Distancing @ Parks, and JumpStart/MyCareersFuture.

B. There is no single best AI solution for customer or public-facing situations because there are many different AI methods, all of which can be used in a range of use cases and domains. Early versions of an application often do not require the most sophisticated AI methods or delegate a great deal to the automated system.

Practitioners must learn to strike a very careful balance between the potential benefits of enhanced and automated accessibility, responsiveness, and personalization and the potential risks of losing public trust, overreaching, opacity, and excessive complexity.

The use of AI for consequential interactions with the public is in its infancy. Leaders must choose projects carefully, in part to persuade the public of the usefulness of AI. For now, we should focus on use cases that are not overly controversial and in which the machine is not solely responsible for high-impact decisions.

Over time, people will develop more trust and comfort with with AI-enabled systems. Meanwhile the solutions already deployed will gradually improve, prompting leaders to expand AI's range of uses and allow AI systems to take

on bigger roles through closely monitored increase of both task automation and AI augmentation.

C. The key to success in each of these examples was a focus on solving a clearly identified problem in order to meet people's important everyday needs. None of these applications were motivated by the desire to use AI for its own sake. AI methods were employed because they were best suited to addressing the users' needs. It is more important to clearly define the use case than to use the most sophisticated AI methods. AI methods and supporting data that produce sufficient results at the outset can be improved over time.

D. Once leaders have a goal in mind, they should determine the broad types of AI functionality required for the initial minimum viable product and for updates over the initial eighteen months. Most projects will require one or more of the following families of functionality: prediction/recommendation, planning/optimization, modeling/simulation, natural language understanding, image or video understanding, or work process automation.

Once they have framed the problem, leaders will find their selection of AI methods influenced by i) data availability, ii) required degree and nature of transparency and auditability, and iii) the risks and consequences associated with AI errors and inaccuracies such as false positives, false negatives, and other output errors.

E. For an entity in the public sector, focusing on addressing the existing and common, even very basic, needs of residents is a good way to get started. It allows the agency to build internal AI capabilities while residents build experience with using AI systems.

The first step to becoming a smart city is to help residents to

do basic, everyday tasks more easily. This foundation will support future progress toward using more sophisticated AI methods and models to offer new services to residents.

Companies, especially those in the early stages of deploying AI tools, should follow this same pathway. Start simple; focus on basic customer needs.

F. The success of the latter five examples are the result of steady, iterative efforts spanning multi-year time periods. The designers needed this extended time to build up datasets, refine the performance of AI models, and integrate their outputs.

Organizational leaders must also plan and support a patient and iterative approach over a period of multiple years, as long as the system demonstrates a steady learning and performance curve at each step.

Example one demonstrates that once the organization has a strong capability base, it will be able to design and deploy targeted solutions very quickly in specific situations.

G. Working simultaneously to establish policies, governance, and technology platforms is essential if an organization hopes to move from an ad hoc approach to each AI project to a portfolio approach, deploying multiple AI solutions throughout the organization.

Being clear about using the commercial cloud, internal data-sharing, ownership and accountability for data quality, data and privacy protection, and use case review and approval is as important as AI technical competency.

Leaders must build internal infrastructure and software platforms that can capture, clean, curate, manage, and distribute data throughout the organization. They must also build the platforms for managing and executing development, security, and operations (DevSecOps) for machine learning.[33] All of these steps are necessary prerequisites for increasing the organization's efficiency in deploying and supporting an expanding portfolio of AI projects.

AI-enabled solutions for customer engagement, or for any application, involve a lot more than AI.

The technical aspects of choosing and implementing AI-models, while critical, are just a small part of creating, deploying and operating an AI-enabled solution. AI-enabled solutions for customer engagement, or for any type of application, involve a lot more than AI.

Building the City's Capacity to be a Lab for AI Deployments

Singapore's strategy is to rapidly test new ideas for AI applications in a city-scale field setting. Singapore can often move more quickly than larger countries to iteratively build a base of common understanding between the realms of technology, infrastructure, management, policy, and governance in order to build better and more complex public sector AI capabilities and services.

Singapore's government is therefore committed to encouraging responsible and innovative field trials and follow-on deployments of existing and emerging AI technologies. Despite the country's small size in terms of geography and population, its approach to field trials in carefully considered use cases with well-orchestrated deployments attracts leading global companies that provide AI solutions as well as AI-focused start-ups.

Uses of AI in the public sector have proven to be good starting points for interagency partnerships and for public-private partnerships. The Singapore government's approach to developing AI tools is a mix of home-building and using the most innovative vendors around the world.

Simpler – meaning no need for major customization – and less sensitive, smaller-impact projects are more receptive to the use of commercial products and open-source library modules to achieve quick wins (e.g., Safe Distancing @ Parks). More complex, highly integrated, and novel projects usually necessitate deeply contextual and custom approaches that require local knowledge and innovations (e.g., the JumpStart platform for the MyCareersFuture job search portal and the MSO chatbot, both internally developed by GovTech).

Some of these complex cases, like the AI learning assistants, the breeze-through system, and the FASTER system, also required joint development partnerships with local and global vendors and research institutions. The FASTER rail management system was achieved through a multi-year partnership with IBM and ST Engineering, a local firm.

Uses of AI in the public sector have proven to be good starting points for interagency partnerships and for public-private partnerships.

The government's learning how to steer relevant regulations, navigating public and societal protections, while supporting R&D, innovation, and learning-by-doing, was especially essential.

To build AI projects, governments must also be prepared to consider policies, regulations, and protections regarding access to public data, because data-driven

machine learning underlies most contemporary AI applications. Not long ago, Singapore was known for its unusually tight grip on government information; even data-sharing within the government was often cumbersome. Today, the situation has changed; access to both government and public data has become much more easy and open, though still with some notable exceptions.

Leadership

The Singapore government's AI journey, including AI for public engagement, has required leaders who have a judicious combination of the following four traits:

- **Progressiveness:** Senior leaders must support a dare-to-try spirit within government units, ministries, and agencies. This attitude is essential for digital and AI innovation for improving customers' engagement and experiences.

- **Pragmatism:** Leaders must make no-nonsense evaluations of whether an AI approach is even warranted. If the quality or quantity of data are insufficient, data-driven machine learning methods will not perform well. When addressing a new problem, choosing the simplest AI solution that best fits the situation will usually be an effective way to get started.

  The initial solution might even be good old-fashioned AI's rule- or logic-based approaches in defiance of the current common wisdom that says to always start with more complicated machine learning methods. Leaders also need hardheaded pragmatism to dismiss proposed projects that do not have a clearly defined purpose, or that are likely to be highly controversial or sensitive.

- **Patience:** Leaders and senior management must realize that, when deploying AI-based

solutions, the system will be far from perfect at the outset. They must take a longer view and be willing to sustain the investments needed for the development team to keep iterating, improving the AI's performance over time as it collects more data and designers refine the AI and build supporting modules. These iterations, which build to the desired level of performance, can take several years.

- **Persistence:** Leaders who plan to introduce a new technology should be prepared to hold steady and persist in the face of internal and external public resistance. Like deploying any new type of technology, launching new types of AI-enabled systems and processes will disrupt established practices.

  The first AI projects of organizations will be opposed by internal and external stakeholders who may have a variety of concerns including privacy, job loss, and fairness. To deploy and adopt AI, leaders must have the conviction and the evidence necessary to responsibly address this inevitable opposition.

  Leaders will also need to be persistent in a way that is closely linked to patience. AI systems of all types require substantial monitoring, adjustment, and maintenance to ensure that their results remain reliable and responsible. Without this persistent support after deployment, the system's performance will drift as conditions change.

  Support teams will also have to address biases and errors in the system, both overt and subtle, that they did not initially detect or understand, but which emerge over time. Leaders must take responsibility for the organizational mechanisms that handle the lifetime requirements of AI systems

> **Leaders must take responsibility for the organizational mechanisms that handle the lifetime requirements of AI systems and products. They cannot just be deployed and forgotten.**

and products. They cannot just be deployed and forgotten.

In summary, the Singapore government's approach to building its AI capability and deploying AI applications across the public sector has been a national-level engagement effort across the whole government and its boundaries, the private sector, and local society.

The result is Singapore's steadily expanding portfolio of AI-enabled applications that have changed both how the government operates and how it engages with the public. So far, residents and businesses have responded positively to the government's digital and AI efforts. Singapore's annual survey on satisfaction with these services reflects that.[34]

On a scale of one to five, with five being 'very satisfied,' overall citizen satisfaction rose from 95 percent satisfied or higher and 73 percent very satisfied in 2016 to 99 percent satisfied or higher and 85 percent very satisfied in 2021. Satisfaction scores from businesses also rose: from 92 percent satisfied or higher, 64 percent very satisfied in 2016 to 98 percent satisfied or higher, 76 percent very satisfied in 2021. While the AI applications in these six examples are only a small part of a much larger government digital effort, the scores show that AI will play an increasingly important and expanding part in Singaporean life for years to come. ■

Author Bio

Steven M. Miller is Professor Emeritus of Information Systems at Singapore Management University. He served as founding dean of SMU's School of Computing and Information Systems for over thirteen years.[35] He also served as SMU's vice provost for research for eight years. He continues to advise several companies and government agencies in Singapore on AI applications and deployments, digital transformation, and organizational learning.

Endnotes

1. Smart Nation Singapore, Smart Nation and Digital Government Office, "National Artificial Intelligence Strategy: Advancing our Smart Nation Journey," November 2019, https://www.smartnation.gov.sg/files/publications/national-ai-strategy.pdf , viewed June 29 2022.

2. Smart Nation Singapore, "National AI Strategy," https://www.smartnation.gov.sg/initiatives/artificial-intelligence, viewed June 29 2022.

3. Special thanks to the staff of the National AI Office within the Smart Nation and Digital Government Office for consultations that led to the selection of this set of examples as good illustrations of applications of AI that emphasize public engagement, and for additional information on these applications. Also, special thanks to staff from GovTech, from the Home Team Science and Technology Agency, and the Land Transportation Agency for their efforts to review the descriptions of these applications and share insights. Any errors related to facts or nuances are the full responsibility of the author.

4. National Parks Board, "Singapore, our City in Nature," https://www.nparks.gov.sg/about-us/city-in-nature, viewed June 29 2022.

5. GovTech, "Safe Distance @ Parks: how AI replaced eye power for crowd counting," October 5 2021, https://www.tech.gov.sg/media/technews/safe-distance-at-nparks, viewed June 29 2022.

6. Rachel Shong, "How we deployed an Artificial Intelligence (AI) solution quickly during Circuit Breaker," MEDIUM, DSAID GovTech channel, June 14 2021, https://medium.com/dsaid-govtech/how-we-deployed-an-artificial-intelligence-ai-solution-quickly-during-circuit-breaker-1c69cf814891, viewed June 29 2022.

7. GovTech, "NParks and SNDGG trial SPOT robot for safe distancing operations at Bishan-Ang Mo Kio Park," May 2020, https://www.tech.gov.sg/media/media-releases/spot-robot-trial-for-safe-distancing-operations, viewed June 29 2022.

8. Shashank Bengali, "Singapore enforces social distancing – with a robot dog," Los Angeles Times, May 13 2020. Photo in article (Figure 3) by Roslan Rahman/AFP via Getty Images, https://www.latimes.com/world-nation/story/2020-05-13/coronavirus-singapores-robot-dog-enforces-social-distancing, viewed 08 August 2022.

9. The Ministry of Home Affairs is also actively working on testing and gradually deploying four legged robots for civil defense work (e.g., response to fires, to hazardous material threats) as well as for police work. See Shermaine Ang, "Four-legged robot can help SCDF respond to hazardous material incidents, assist police with patrols," Straits Times, January 20 2022, https://www.straitstimes.com/singapore/trials-for-new-hazmat-function-of-robotic-dog-rover-x-to-be-completed-by-mid-year, viewed June 29 2022.

10. National Parks Board, "Safe Distance @ Parks," https://safedistparks.nparks.gov.sg/, viewed June 29 2022.

11. Committee on the Future Economy, "Report of the Committee on the Future Economy," February 2017, https://www.mti.gov.sg/Resources/publications/Report-of-the-Committee-on-the-Future-Economy, viewed June 29 2022. Also see Grace Ho, "Industry Transformation Maps to be refreshed, strengthened over next 5 years: Deputy Prime Minister Heng," Straits Times, April 30 2021, https://www.straitstimes.com/singapore/politics/industry-transformation-maps-to-be-refreshed-and-strengthened-over-the-next-five, viewed June 29 2022.

12. SkillsFuture Singapore, "About Skills Future," https://www.skillsfuture.gov.sg/AboutSkillsFuture, viewed June 29 2022.

13. Workforce Singapore, "MyCareersFuture" portal, https://www.mycareersfuture.gov.sg/, viewed June 29 2022.

14. Workforce Singapore, "Frequently Asked Questions For MyCareersFuture," https://www.mycareersfuture.gov.sg/docs/mycareersfuture_sg_user_faqs.pdf, viewed June 29 2022.

15. Joanna Seow, "New jobs portal MyCareersFuture.sg uses technology to better match jobseekers and employers," Straits Times, April 17 2018, https://www.straitstimes.com/singapore/manpower/new-jobs-portal-uses-technology-to-better-match-jobseekers-and-employers, viewed June 29 2022.

16. Skills Future SG and Workforce Singapore, "MyCareersFuture now open for employers; new features to quicken matching process between employers and jobseekers," September 2 2019, https://www.ssg-wsg.gov.sg/news-and-announcements/2-Sep-2019.html, viewed June 29 2022.

17. Clarence Ng, Elizabeth Lim, Tay Mingfang, "Using Recommender Systems to Improve Job Search with the JumpStart Platform," MEDIUM, DSAID GovTech channel, December 3 2021, https://medium.com/dsaid-govtech/using-recommender-systems-to-improve-job-search-with-the-jumpstart-platform-cbfbb96913ca, June 29 2022.

18. Fred Teo, Clarence Ng, Elizabeth Lim, and Lua Jiong Wei, "Generating Job Recommendations for Jobseekers on MyCareersFuture," MEDIUM, DSAID GovTech channel, December 6 2021, https://medium.com/dsaid-govtech/generating-job-recommendations-for-jobseekers-on-mycareersfuture-699c82cf6463, viewed June 29 2022.

19. Steven M. Miller and Thomas H. Davenport, "A Smarter Way to Manage Mass Transit in a Smart City: Rail Network Management at Singapore's Land Transport Authority," AI Singapore website, May 27 2021, https://aisingapore.org/2021/05/a-smarter-way-to-manage-mass-transit-in-a-smart-city-rail-network-management-at-singapores-land-transport-authority/, viewed June 29 2022.

20. Municipal Service Office, "Who We Are," https://www.mnd.gov.sg/mso/who-we-are/vision-and-mission, viewed June 29 2022.

21. Smart Nation Singapore, "OneService App," https://www.smartnation.gov.sg/initiatives/urban-living/oneservice-app, viewed June 29 2022.

22. Smart Nation Singapore, "New AI-powered OneService Chatbot for easy reporting of municipal issues," July 15 2021, https://www.smartnation.gov.sg/media-hub/press-releases/new-ai-powered-oneservice, viewed June 29 2022.

23. GovTech, "VICA - Virtual Intelligent Chat Assistant," https://www.tech.gov.sg/products-and-services/vica/, viewed 29 June 2022.

24. GovTech, Singapore Government Developer Portal, "Virtual Intelligent Chat Assistant (VICA) – A Natural Language Processing (NLP) Chatbot For Citizens," https://www.developer.tech.gov.sg/products/categories/platform/virtual-intelligent-chat-assistant/overview.html, viewed 29 June 2022.

25. Ministry of Education, "Singapore Student Learning Space (SLS)," https://www.moe.gov.sg/education-in-sg/student-learning-space, viewed 29 June 2022.

26. Ministry of Education, Parliamentary reply to Member of Parliament question on "Adaptive Learning," July 5 2021, https://www.moe.gov.sg/news/parliamentary-replies/20210705-adaptive-learning, viewed 29 June 2022.

27. Ministry of Education, "Advance at Your Own Pace: More personalised learning with the use of Adaptive Learning enhanced by Artificial Intelligence!," Contact: The Teachers' Digest, Issue 36/37, March 2020, https://www.moe.gov.sg/about-us/publications/contact-teachers-digest/our-chance-to-shape-the-future, viewed 29 June 2022.

28. Ng Wei Kai, "Upper primary, secondary students to get instant feedback on English assignments with new AI system," Straits Times, 18 August 2021. https://www.straitstimes.com/singapore/parenting-education/moe-to-roll-out-ai-enabled-marking-system-for-english-language-writing, viewed 29 June 2022.

29. Ministry of Education, "Parliamentary reply to Member of Parliament question on Automated Marking System" 03 November 2021, https://www.moe.gov.sg/news/parliamentary-replies/20211103-automated-marking-system, viewed 29 June 2022.

30. Immigration & Checkpoints Authority, "Use Of Iris And Facial Biometrics As The Primary Biometric Identifiers For Immigration Clearance At All Checkpoints," October 28 2020, https://www.ica.gov.sg/news-and-publications/newsroom/media-release/use-of-iris-and-facial-biometrics-as-the-primary-biometric-identifiers-for-immigration-clearance-at-all-checkpoints, viewed 29 June 2022.

31. Home Team Science and Technology Agency (HTX), "Biometrics and Profiling," https://www.htx.gov.sg/expertise/our-expertise/biometrics-profiling, viewed 29 June 2022.

32. GovTech, "Singapore Digital Government Journey," https://www.tech.gov.sg/singapore-digital-government-journey/. See the four video interview episodes under "Our Digital Government Story." In these video interview episodes, the Government Chief Digital Technology Officer explains how they built up the capability to do a wide range of digital transformation projects – including AI projects – and the key principles they followed and policies they developed in order to make progress with digitalization and AI efforts at scale and at speed over the 2014 to 2022 timeframe.

33. Luiz Antunes, "Aligning DevSecOps and Machine Learning," Software Engineering Institute, Carnegie Mellon University, May 3 2021, https://insights.sei.cmu.edu/blog/aligning-devsecops-and-machine-learning/, viewed July 31 2022.

34. GovTech, Our Statistics, "Annual Survey on Satisfaction with Government Digital Services For 2021, (Citizens) and (Business)," https://www.tech.gov.sg/who-we-are/our-statistics/, viewed 06 September 2022.

35. https://scis.smu.edu.sg/

Marketing Artificial Intelligence: AI, Marketing, and the Future of Business
by Paul Roetzer and Mike Kaput

Michael Diamond

New York University

Book synopsis:

Marketers expect artificial intelligence (AI) to have trillions of dollars of impact on businesses and the economy, yet many struggle to understand what it is and how to apply it to their efforts. But businesses do not need to begin by mounting a gargantuan AI initiative. Instead, the authors argue, deploying a little bit of AI can go a long way toward increasing a company's productivity, efficiency, and performance. They describe AI's current potential and offer a glimpse into a future in which marketers and machines unite to run personalized and complex campaigns with great simplicity.

Paul Roetzer and Mike Kaput's book, *Marketing Artificial Intelligence*, does an excellent job in three important areas:

1. It builds the case for businesses to adopt AI in their strategy, technology, and talent;
2. It explains the foundational concepts of AI and its subspecies – machine learning, deep learning, conversational AI, and more;
3. It establishes foundational frameworks for businesses to use in analyzing opportunities, assessing vendors, and internally reflecting on their own talent, capabilities, and organizational maturity.

Along the way, the book provides a series of well-structured, pragmatic explorations of AI's impact on brand building, campaign development, and experience design, as well as strategy making, business planning, and innovation.

> Moving forward, a company's capacity to wield AI to build richer engagements with customers will be one of the strongest single markers of competitive advantage.

Marketer plus machine
The book's foundation is the idea that "AI enhances human knowledge and capabilities … [so that] the future is marketer plus machine." The authors highlight how machine-based intelligence is driving an exponential rate of change in the marketing field.

Much of the authors' insight about how AI enhances our marketing ability applies to how we come to understand consumer engagement (CE). Central to CE is the desire of humans for greater personalization and control in a complex world of choices. Using AI in consumer engagement allows brands to give agency back to the customer.

Roetzer and Kaput's central argument is that, moving forward, a company's capacity to wield AI to build richer engagements with customers will be one of the strongest single markers of competitive advantage.

Human-centered and data-driven
Roetzer and Kaput's book reflects the idea that deepening customer engagement will always be both human-centered and data-driven. The authors argue that the current state of AI can more concretely fulfill the ambitions of anticipatory design and the intuitive Internet, which have informed discussions of customer engagement over the past two decades.

The authors unpack how AI delivers deeper insights about consumer experience and behavior than were previously available. They argue that given context, prior activity, and design elements, AI can accurately predict consumer behaviors, which can then improve the customers' experience.

The book provides much-needed frameworks and benchmarking tools which marketers can use to gain a deeper understanding of customer engagement. These instruments provide a strong, emerging set of criteria, testable in specific, individual cases and include: the Marketer-to-Machine Scale, which classifies five levels of automation so users can better understand how strongly a vendor's offerings might affect their business; and the 5Ps of Marketing AI, a framework which we can use to analyze the current technological landscape.

As for Roetzer and Kaput's 5Ps, planning, production, personalization, promotion and performance, three are particularly germane to the development of customer engagement: planning, personalization, and performance.

Planning

Building intelligent strategies calls for constructing accurate "personas based on needs, goals, intent and behavior," among other traits. Authentic and responsive customer engagement requires that advertisers understand how customers connect with market strategies, evaluate how these connections create value, and respond to customers as they create connections.

Personalization

Personalization, "powering intelligent consumer experiences," ensures that AI can drive both a more targeted set of incentives and overtures to engage consumers and more relevant content and conversations, through which it continues to learn and improve.

Performance

Performance, "turning data into intelligence," uses AI to evaluate the return on investment (ROI) of customer engagement efforts. AI helps to determine how to reallocate investments across initiatives in response to new behaviors, such as increases in customer loyalty and satisfaction.

The broader challenge for practitioners applying artificial intelligence to customer engagement, however, is that CE is "a multidimensional concept" with "relevant cognitive, emotional and/or behavioral dimensions."[1] AI's ability to enhance and influence these interactions changes as AI approaches human sentience limits. The advancement of AI into the domain of human sentience, it should be noted, brings up ethical risks to companies' reputations.

Marketing Artificial Intelligence is a practical, well-researched, and applied introduction to how marketers can explore the new opportunities and directions which AI has created to deepen and enhance customer engagement. ■

Author Bio

Michael Diamond is a clinical assistant professor of integrated marketing and communications at NYU's School of Professional Studies. He was previously the acting CMO of Time Warner Cable and a strategy consultant for media, entertainment, and technology companies with Booz Allen Hamilton.

Endnote

1. Brodie, Roderick J, Linda D. Hollebeek, Biljana Jurić, & Anna Ilić. "Customer Engagement: Conceptual Domain, Fundamental Propositions, and Implications for Research," Journal of Service Research 14, no. 4 (2011), 252–271.

Book Review: *Quantum Marketing*
by RAJA RAJAMANNAR

Nick Primola, Association of National Advertisers

Quantum Marketing[1] provides a foundation for anyone seeking to understand the current wave of tech-driven innovations that are transforming customer engagement. The author, Raja Rajamannar, is the Chief Marketing and Communications Officer of Mastercard and a globally recognized marketing expert, innovator, transformational business leader, and world-class brand builder. In *Quantum Marketing*, Rajamannar uses his first-hand experience and personal insights to bring AI to life in a clear and practical way.

Quantum Marketing refers to how the world is changing at such an unprecedented rate that classical marketing theories, strategies, and practices are failing to meet consumers' and marketers' needs. *Quantum Marketing* provides a new framework to help us understand this new world.

Understanding the Role of Artificial Intelligence in this Moment

Prior to the explosive emergence of generative AI, Rajamannar stated, "Nothing will change marketing more than AI." Marketers can use it to learn about consumers deeply, so they can hyper-personalize and optimize programs in real time. It is essential for marketers to learn to use AI to keep up with industry and consumer expectations. *Quantum Marketing* helps the reader to understand how AI will transform customer engagement in several ways:

Quantum Marketing sets the benchmark for real-time marketing

Using AI, companies can collect customer insights in real-time, so they can respond with near-zero lag. Marketers can measure the impact of these responses instantly and optimize their efforts immediately: that is quantum marketing. AI lets marketers monitor all stages of the marketing life cycle, make sense of them, and act to yield the most effective outcomes.

AI will help marketers to drive growth through near-perfect relevance

A common goal for all marketers is creating advertising that is seamless, unintrusive, and not annoying. This goal is part of an ongoing give and take between consumers and marketers. For relevant and appealing marketing, consumers must be willing to accept the presence of sensors everywhere – from watches to shoes, cars, golf clubs, thermostats, and more. These sensors capture data that AI machines analyze to create amazing patterns and insights that marketers can respond to in real time. Being able to create value for consumers at optimal moments of their daily lives gives these marketers a significant advantage.

AI is raising the bar for customer engagement

AI will enable advertisers to create consumer experiences that are relevant and seamless. Marketers work in an environment with ever-increasing expectations as companies leverage new technology and innovations. Rajamannar states that marketers need to start experimenting now with the many options already available, from off-the-shelf solutions to those from companies like Amazon and Google.

A Warning for Marketers – You Are Responsible for Maintaining Societal Trust

Consumers are increasingly uncertain about the veracity of what they see and hear. As artificial intelligence becomes ever more prevalent, it will aggravate this situation, introducing fake photos, videos, and voices, fake just about anything. This danger is already affecting politics, societal norms, and professional integrity. That AI can be used to create these fakes is a reality marketers must face. The proper use of AI – whether machine-to-person or machine-to-machine – must be undertaken thoughtfully. Marketers must be aware of the potentially devastating effects of using AI irresponsibly.

In summary, *Quantum Marketing* describes the context and environment from which marketers will reach new horizons and the expectations of consumers will be elevated to unprecedented heights. Those who effectively leverage artificial intelligence will enjoy a clear competitive advantage. AI is essential for marketers to engage consumers on their terms while keeping pace with an ever-rising tide of expectations. ■

Author Bio

Nick Primola is Group EVP and head of the CMO Practice of the Association of National Advertisers. In partnership with Cannes LIONS, he runs the Global CMO Growth Council. Nick was previously senior vice president of corporate marketing at Citizens Bank and is a veteran of the world-class agencies Ogilvy and Young & Rubicam.

Endnote

1. https://quantummarketing.com/

3 Visions of the Future of AI for Customer Engagement: 2027 Scenarios

©shutterstock.com/Dan Collier

Jerry Wind
The Wharton School, University of Pennsylvania

Mukul Pandya
The Wharton School, University of Pennsylvania

Margherita Pagani
SKEMA Business School

Jerry Dischler
Google

Artificial intelligence (AI) is making deeper inroads into every aspect of business and society every day. By 2027, how will AI transform the future of customer engagement? *Management and Business Review* and the Association of National Advertisers' (ANA) Global Chief Marketing Officer (CMO) Growth Council recently organized a forum to discuss three possible scenarios for the future – optimistic, pessimistic, and realistic.

The goal was to develop consensus on the most likely ways in which AI will transform customer engagement in the next few years. Most importantly, participants focused on what steps to take today to prepare for 2027. These

insights are contained in this special issue of *Management and Business Review* on customer engagement.

The forum was moderated by Professor Jerry Wind, the Lauder Professor Emeritus and professor of marketing at the Wharton School of the University of Pennsylvania and editor of this special MBR issue on AI for customer engagement, along with Nick Primola, ANA executive vice president and director of its CMO Growth Council. Samantha Stetson, vice president client council and industry trade relations at Meta, presented opening remarks.

The event aspired to "agree on a most realistic future-state so that we could help business executives plan their roadmap to take action," Primola said. "And – importantly – for the industry to work together to ensure that we end up closer to the more optimistic state versus the most pessimistic."

When work began on this special issue, generative AI did not have the momentum it developed in the beginning of 2023. It is clear now that generative AI holds tremendous promise. This fast-developing wave was further energized in March 2023, when OpenAI launched GPT-4, its latest large language model. The announcement on Open AI's website said that "GPT-4 is OpenAI's most advanced system, producing safer and most useful responses." According to OpenAI, the new software "can solve difficult problems with greater accuracy, thanks to its broader general knowledge and problem solving abilities."

Early experiments indicate that GPT-4 could have a profound impact upon every industry. *The New York Times* columnist Thomas L. Friedman, after seeing a demonstration of GPT-4, quoted Arthur C. Clarke, who wrote that "any sufficiently advanced technology is indistinguishable from magic." Friedman also quoted Craig Mundie, the former chief research

and strategy officer for Microsoft, who said, "This is going to change everything about how we do everything…[I]t represents mankind's greatest invention to date."

Other observers have been equally optimistic. Shelly Palmer, a professor of advanced media in residence at the Newhouse School of Public Communications and CEO of The Palmer Group, says that with GPT-4, generative AI comes close to "sparks of artificial general intelligence," which means that it has capabilities that are at or above human level. Palmer believes that "by assisting its users with a wide variety of tasks," generative AI tools such as GPT-4, if used well, "will be a force multiplier for productivity across every enterprise."

All this was still in the future in 2022 when the workshop considered the three scenarios. Yet these developments in generative AI do not change the importance of the conclusions we drew at the workshop. Instead they suggest the urgency of understanding the implications of current developments on AI and customer engagement and the need for constant experimentation and decisive action.

The optimistic scenario

The increasing deployment of AI should bring about several benefits, one of which is its ability to parse big data to get at the good data. "At present, 10 percent of the data that companies gather is useful, and 90 percent is junk," one participant said. "AI will help to convert the junk into useful data."

Auto insurer Metromile, for example, introduced a pay-as-you-go model based on granular customer data of actual driving behavior. Customers who drive fewer miles along less congested roads could pay lower premiums than those who drive longer

distances along crowded highways.

Some participants said AI permits companies to take a 360-degree view of their data, which should increase customer engagement. Companies will also be able to respond faster to outlier events like supply shortages and support end-to-end value chains.

Another optimistic view was that AI would enhance creativity in two ways. The first through AI-powered tools that boost human

Table 1: Optimistic scenario
Artificial intelligence is expected to
• Raise customer engagement and satisfaction through increased personalization
• Raise customer satisfaction by offering faster, cheaper services
• Improve the usefulness of data
• Bring customers closer to the board room
• Enhance the capabilities of augmented reality
• Provide a 360-degree view of data and make it more accessible
• Support end-to-end value chains
• Expand human creativity
• Optimize change management
• Foster innovation
• Automate repetitive tasks, freeing humans to focus on higher-value activities
• Enhance human interaction
• Enhance trust by deftly identifying fakes
• Correct information imbalances
• Foster better partnerships and segmentation
• Encourage self-governance and ethical AI
• Improve 24/7 customer service
• Be tailored to support professionals
• Improve quality control and checks and balances
• Improve efficiency and optimization
• Enable a broad range of access via education
• Improve media buying
• Increase transparency

creativity, making it possible to produce artistic works such as images or paintings with simple text prompts.

The second is by eliminating the drudgery of tasks that stifle creativity. Photographers who are freed from monotonous and repetitive chores like image cropping could spend more of their time on creative tasks like composing images. The integration of AI with other advanced technologies, such as augmented reality (AR), should spur creativity.

Some participants said AI would affect customer engagement in the next few years by allowing companies to proactively understand and anticipate the needs of customers. "Companies today are looking at AI as a means of cost efficiency, process efficiency, which is all well and good, but the end game should be meeting customers where they are," a participant said.

Some noted that AI will be important in building trust between companies and their customers. They pointed out that as companies produce more content, AI will make it possible for them to identify and eliminate fraudulent content such as deep fakes – even as AI can also create fake content – with the result that content will become more credible and trustworthy.

Eliminating fake content from the Internet would be impossible without the speed and scale of AI. "AI will increase transparency," one participant noted. "Can AI verify facts?" another asked. If it can, it will increase trust, which is critical to customer engagement.

AI would also improve human interaction, participants noted. It not only enables the creation of better content faster through natural language generation, but allows for much more precise delivery to the right audience. AI will make content producers more efficient as well as "support the expression of diverse viewpoints," a participant said.

The pessimistic view

For almost every positive aspect of AI, forum participants who took the pessimistic view saw the corresponding downside.

Take, for example, the view that AI will increase the usefulness of data; it could be a double-edged sword. Consider auto insurance again. If AI can crunch sensor data about each trip to better assess risk and potentially lower premiums, perhaps it can also discover what the car's occupants are doing. Not only is this an invasion of privacy, but "if your kids are watching TV in the car, you might have to pay more," said a participant.

Moreover, bringing fresh streams of data into decisions about pricing insurance at variable rates for different customers does not mean that these decisions will be, or be perceived to be, fair. "If I am an Uber driver who doesn't make much money, I might be paying more in insurance costs than what I earn driving for Uber," said a participant. "Taxing the poor this way can expose a lot of inequities. A rich person who drives little may have to pay less than a poor person who has no choice but to drive for work."

The view that AI can increase trust was also disputed by some. One participant said that AI can undermine trust by encouraging decisions that are entirely transactional rather than human or empathetic. "My view is that the deployment of AI, either intentional or unintentional, undermines trust in systems. The number one way that AI is used for behavior modification is in e-commerce, in getting people to buy things. If it turns out that they are being prompted to buy things that they do not necessarily need, people may perceive AI as being manipulative rather than helpful. AI uses math to make predictions, but that does not mean that the predictions will be ethical."

The ethics of AI will continue to pose an ongoing challenge, according to the forum participants. Although neural networks and deep learn-

Table 2: Pessimistic scenario

Artificial intelligence could lead to

- Worsening of the silo effect, leading to social polarization if paired with social media
- Overcrowding of the search engine optimization space if everyone uses AI for it
- Loss of privacy
- Exclusion of certain groups
- A loss of the human factor in increasingly transactional decisions
- A fuzzy boundary between AI and human effort
- A need for guardrails even if AI is an incredible tool
- Challenges in regulating and setting governance standards
- A view that it is more manipulative than helpful in its selling efforts
- Increased societal risk through manipulation of information, especially video
- Increased mistrust
- Loss of creativity
- Biased outcomes due to being trained on biased data
- Worsening of the wealth gap and social disharmony
- Pollution of content

ing have had an enormous impact in many areas, if the training data is inherently biased for historical reasons, the predictions and recommendations of AI will reflect that bias. This is especially true when AI is used for predictions in areas such as judicial programs, where future potential offenders are often predicted to belong to certain racial groups based on biased historical data.

The pessimistic scenario also regarded the impact of AI upon human interaction and communication in darker terms than the optimistic one. When AI algorithms or bots are paired with social media platforms, they seem to aggravate social fragmentation and division rather

than the reverse. As one participant noted: "All AI will become a technology that social media will exploit to further polarize society."

Most realistic scenario for 2027

According to participants at the forum, the most likely scenario for 2027 will be a blend of the optimistic and pessimistic views. For example, AI will continue to make deep inroads in many industries, with more enhancements to come. "AI and humans will continue to support each other in improving the overall experiences that customers are getting and creating new channels of engagement for customers," a participant noted.

Another participant noted that "we will see more AI-human collaboration." However, "we will see fewer cases where marketers let the AI run wild, and more cases where AI is producing things within the confines of marketers, which allows teams to get more done with teams of the same or smaller size. Another thing is, (the AI wild west) will get worse before it gets better in terms of companies claiming to be AI when they are really not."

AI will supercharge and democratize creativity. "The ability of AI to synthesize content will move to a capability that the world has never had before, on still images, video images, and written content," one participant said. "That will have massive implications. That capability won't just be available for the biggest and richest companies; it will be widely available."

In addition, by 2027 companies will substantially increase their investment in educating employees about AI. While today the public is at a very early stage of understanding AI, with many who have heard about AI but do not know exactly what it means, people will be motivated to learn more about it as more industries start using AI extensively. "This will drive investments in AI education,"

a participant said. "This is a realistic scenario."

The years to 2027 will also see increasing integration of AI with the metaverse, often described as the next iteration of the Internet, marked by increasing use of augmented reality and virtual reality. "Reinventing AI for the metaverse is something that currently is not complete," said a participant. It will move much further along in the next few years.

Some participants noted that by 2027, the most prevalent use of AI in business systems will be for optimization. AI will help organizations get things done faster and more effectively. "The other area where we think AI is going to be effective is in what we call mission-safe applications," a participant noted. "Netflix is a mission-safe application. You can try an experiment on Netflix and see if the viewer likes it. Areas like e-commerce and entertainment are not dangerous, and AI will proliferate there."

What is to be done today?

In addition to discussing these three scenarios, the forum focused on the importance of users considering what will position them well for the future, while examining the implications of AI on customer engagement today. "The important thing is that every one of us should try to consider what these ideas mean. How can we actually leverage AI? AI is going to be with us forever. How will we increase its impact?" said Professor Wind. "We have to think about this going beyond the current utilization. Think about the innovative new insights we can gain as a result of this amazing technology; and think about human-AI interaction as the new unit of analysis in each area with which we are dealing."

Ten Scenarios for 2041 - from Kai-Fu Lee and Chen Qiufan

While the forum focused on the period up to 2027, it may also be useful to look further out to see how AI

Table 3: Action items by 2027
Companies should
▪ Experiment on how AI can increase productivity
▪ Find ways to enhance human-AI collaboration
▪ Improve regulation to dispel inaccurate claims about AI
▪ Leverage AI to improve efficiency and data optimization in various fields
▪ Deploy AI to influence environmental, social and governance (ESG) concerns and ethical marketing
▪ Focus on human-centric AI, with people defining virtues and values
▪ Focus on empathy
▪ Avoid letting AI turn the Internet into a 'splinternet'
▪ Recognize that AI will touch more and more areas, outpacing customer understanding
▪ Demystify AI by educating all customers
▪ Prioritize ideas that enhance customer engagement
▪ Meet customers where they are, based on their needs and interests
▪ Understand the need for clarity about black-box AI processes
▪ Move from AI as a high-tech asset to broad access
▪ Understand and act upon the relationship between AI and creativity

might impact customer engagement, business, and society over the next twenty years. To do this, we considered the book, *AI 2041: Ten Visions for Our Future*, by Kai-Fu Lee and Chen Qiufan.

Kai-Fu Lee is the CEO of Sinovation Ventures and the author of *AI Superpowers*, a book that deals with the emerging AI rivalry between the U.S. and China. A former president of Google China, he is now cochair of the Artificial Intelligence Council at the World Economic Forum. Chen Qiufan, also known as Stanley Chan, is an author and translator, and founder

of the content studio Thema Mundi. Chen was Kai-Fu Lee's colleague at Google. He is the president of the World Chinese Science Fiction Association and the author of *Waste Tide*, published in English in 2019.

The book *AI 2041* is an unusual blend of science fiction and analysis. It contains ten scenarios, or visions, of what is likely to become reality over the next twenty years. These visions take the form of ten science fiction stories written by Chen, each set in a different part of the world. Each story is followed by commentary from Lee. The combination of compelling sci-fi with real-life analysis makes the book fascinating reading.

Of the ten stories, the first seven cover how different industries might apply AI technology. The rest focus on social and geopolitical issues raised by AI.

1. **The Golden Elephant**: Set in Mumbai, India, the story introduces the reader to the basics of AI and offers a sense of its strengths and weaknesses. It asks what risks might emerge when one company possesses a tremendous amount of user data, considering the applications of big data and their privacy implications. "Deep learning makes it possible to customize content for individual users, showing them content that is most likely to appeal to each of them," Lee writes. "This produces greater customer engagement and clicks and purchases than the static, traditional websites were able to do. As companies collect more data about the users, they make more money."

2. **Gods Behind the Masks**: The story revolves around a Nigerian video producer who tries to make undetectable deep fakes with dangerous consequences. Masks become surveillance devices. As Lee notes in his comments, "Computer vision teaches computers to 'see' and recent breakthroughs allow AI to do so like never before." He explores impending breakthroughs in "computer vision, biometrics, and AI security" that make it possible to create deep fakes and other forms of cybercriminal activity.

3. **Twin Sparrows**: This story explores the future of AI education embodied by smart AI teachers "camouflaged as virtual cartoon-like friends." This story is based on natural language processing, which according to the authors is "poised for a meteoric rise in the next decade." Other themes that Lee explores in his commentary are self-supervised training, GPT-3 (or any of the more advanced large language models), artificial general intelligence as well as AI education.

4. **Contactless Love**: Set in China, this story deals with AI in health care. It "explores some of the questions raised by the arrival of a globe-altering pandemic, including its stresses, and also how COVID-19 has accelerated trends such as drug discovery, precision medicine, and robotic surgery," according to Lee.

5. **My Haunting Idol**: This story and commentary depict the future of entertainment, where "games become immersive and the boundary between the real and the virtual is blurred." Set in Tokyo, this story is about how a fan investigates the death of her idol who is brought back to life using virtual reality. In his commentary, Lee discusses the blending of AI with virtual reality, augmented reality, mixed reality, and the brain to computer interface, as well as ethical and societal issues. As the metaverse evolves further and is integrated with AI, this has profound implications for the future.

6. **The Holy Driver**: The story is set in Sri Lanka and imagines a society that is transitioning from human to autonomous driving, powered by AI. In his comments, Lee discusses autonomous vehicles, smart cities, and their ethical and social implications.

7. **Quantum Genocide**: One of the most powerful stories in the book, the story is about a European computer scientist who becomes unhinged and sets upon a course of destruction straight out of a James Bond film. It is set in Iceland. In his comments, Lee describes how quantum computing could turbocharge AI over the next couple of decades.

   "AI-enabled autonomous weapons could become an existential threat to humankind," he writes. Lee notes that in the past, nuclear weapons were seen as a powerful threat to the world, but they are expensive and only a few countries have them. In contrast, the danger of autonomous weapons is that their cost is low. For about $1,000, anyone can equip a drone with facial recognition technology, GPS, and dynamite, creating an assassination machine. "Terrorists do not have to sacrifice their lives if they use these kinds of suicide bombs." Lee believes that countries are not taking this emerging threat seriously enough.

8. **The Job Savior**: Set in the U.S., the story explores what will happen to human jobs as more industries adopt AI. "Technologies such as robotics and robotic process automation will evolve and take over tasks for white- and blue-collar workers," Lee notes. His commentary covers job losses caused by AI and potential solutions such as universal basic income.

9. **Isle of Happiness**: In his commentary, Lee wonders, while

AI may make us efficient and wealthy, can it make us happy? The context for his question is the book's ninth story about a Middle Eastern monarch who wants to use AI as an "elixir for contentment." In addition to questioning whether AI can be used to create happiness, Lee addresses how regulation and privacy factor into the equation.

10. Dreaming of Plenitude: Lee notes in his commentary that "AI will drive down the cost of almost all goods, most of which will be produced for next to nothing." This story is set in Australia and explores abundance, new economic models, and the future of money.

By combining science fiction with factual analysis, the ten powerful stories in this book help readers to understand the likely impact of AI on customer engagement, business, and society well beyond 2027.

Concluding note from the editors:

The 2027 scenarios – optimistic, pessimistic, and most likely – as well as the actions that companies need to take today to prepare for a world in which AI will drive customer engagement and much else, were developed over two and a half hours in an interactive, online-offline forum with more than seventy experts from industry and academia.

As we explore further and deeper into likely future scenarios with larger and more diverse groups, they will enrich these conversations. To balance the session's outcome with a different and longer-term perspective, we have included a summary of *AI 2041*.

The current hype about generative AI does not reduce the importance of the recommendations we developed at the workshop. On the contrary, generative AI highlights the need to act now. The current environment will turbocharge involvement of AI in the way we live, work, and play. That will make understanding and preparing for the future of AI in customer engagement more essential than ever.

Our hope is that the readers of this special issue, keen to dive into innovative applications, thought leadership, and insights about AI for customer engagement, will treat these ideas not just as interesting reading, but as a basis for action. We hope you will use them as a starting point to build your own scenarios, construct your own experiments, and see how you can use AI to build greater customer engagement while sidestepping its pitfalls.

Designing and conducting experiments with AI to enhance customer and human engagement is the best way to get the most from this issue and to learn valuable lessons about this important topic. We hope we have enriched your journey into using AI to enhance your customers' engagement. ■

Author Bios

Yoram (Jerry) Wind is the Lauder Professor Emeritus and Professor of Marketing at Wharton. He founded and long directed the SEI Center for Advanced Studies in Management. He has published more than 300 articles and contributions, and 30 books. An AMA Fellow, Legend in Marketing, and consultant to companies and nonprofits, he holds an honorary doctorate from Reichman University. His research concerns the impact of AI on marketing-driven business strategy, creativity, and innovation.

Mukul Pandya is the founding editor in chief and executive director of Knowledge@Wharton, the online research and business journal of the Wharton School. After retiring from K@W, Mr. Pandya was a senior fellow with Wharton Customer Analytics and AI for Business. A four-time award winner for investigative journalism, Mr. Pandya has published articles in *The New York Times, The Wall Street Journal, the Economist, Time, The Philadelphia Inquirer,* and more. He has written or coauthored four books.

Margherita Pagani, Ph.D., HDR, is a Professor of Digital and Artificial Intelligence in Marketing, and Director of SKEMA Research Centre for Artificial Intelligence at SKEMA Business School (France). She is an advisor for the European Economic and Social Committee (EESC), a research affiliate at the Berkeley APEC Study Center (Basc,) and associate editor for Micro & Macro Marketing. Her research focuses on the impact of AI on consumer behavior and digital ecosystems.

Jerry Dischler is a Vice President and General Manager at Google, overseeing Google's advertising products. A 15+ year Ads veteran, Jerry led Google beyond plain text ads to deeper experiences, developing Ad Sitelinks and Product Listing Ads and helping to define the future of advertising at Google. Before Google, Jerry held a variety of engineering and leadership positions. He earned his bachelor's degree from the University of Chicago and his MBA from The Wharton School.

Executive Summaries

Section I: Personalization
The Mastercard Digital Marketing Engine™: Using AI to Spot Micro Trends for Effective Customer Engagement Raja Rajamannar, Mastercard

People are now bombarded with around 10,000 advertising messages daily, interrupting and impeding their online activities. It's clear that the traditional ways of advertising are no longer viable. Brand marketers must develop a new mindset to connect with consumers and drive return on investment (ROI). The single most powerful enabler and disruptor of marketing is now artificial intelligence (AI). It is an incomparable and essential tool, allowing marketers to engage with consumers meaningfully, authentically, and in real time. Marketers at Mastercard have risen to this challenge and created the Mastercard Digital Engine™. The engine matches consumers' passions and desires with micro trends, instantly delivering a personalized message suited to the context. We are just beginning to scratch the surface of AI's vast potential, but already it is fueling a higher standard of marketing. For marketers at Mastercard, AI is a gamechanger.

The Alibaba Challenge: How to Effectively Engage with a Billion Customers Yitong Wang, Alibaba; Ofer Mintz, University of Technology Sydney and Tel Aviv University; Depin Chen, Alibaba; Kehan Chen, Alibaba

Alibaba uses AI to engage with several million customers daily. Without AI, engaging with this throng of customers, both proactively and reactively, would be functionally impossible. Alibaba therefore uses five AI-based customer service chatbots, each with its own purpose, technical challenges, and abilities. By using AI-based service that provides automated and interactive customer service dispute resolution, Alibaba has achieved a 25 percent increase in customer satisfaction.

AI in Personalized Product Recommendations
Kartik Hosanagar, The Wharton School, University of Pennsylvania; Dokyun Lee, Questrom School of Business, Boston University

Consumers today are presented with a wealth of product offerings in digital marketplaces and storefronts. Managers must consider how consumers will find products of interest among seemingly endless alternatives. Recommender systems combine data drawn from clickstreams, purchases, product ratings, user profiles, and social networks to help consumers learn about new products and select ones. For firms, recommender systems can turn browsers into buyers, cross-sell products, and increase customer loyalty. In recent years, designers have applied machine learning (ML) to create hybrid systems. This method shows great promise in terms of using new kinds of unstructured data to generate personalized recommendations. However, managers should focus not only on short-term customer engagement metrics, but also on long-term customer engagement and the societal impacts of large-scale personalization. It is not necessary to sacrifice the accuracy of recommendations in order to increase their diversity. We can now circumvent the filter bubble effect without causing a drop in customer engagement.

Hyper-Personalization for Customer Engagement with Artificial Intelligence Thomas H. Davenport, Babson College and Oxford Business School

Marketers have, for some time, been familiar with personalization rooted in customer attributes and behavior, and artificial intelligence is making personalization increasingly effective. AI-based hyper-personalization employs sophisticated methods and more data

than previous methods to produce greater precision. AI-fueled personalization increases customer engagement by suggesting products and services that are particularly suited to the specific needs and desires of a given consumer. However, marketers should be aware that the public's growing concerns about data privacy are now generating an increasing backlash against personalization.

Saving Lives with AI: Lessons in Personalization and Engagement Rex Briggs, MMA Global; Stefanie Friedhoff, Brown University; Erik Lundberg, ArtsAI

Fewer than 1 percent of marketers engaged in digital display, audio, and video advertising currently use artificial intelligence (AI) to personalize their ads. Yet this application of AI promises to increase the engagement of customers and improve results for businesses. When charged with teaching people about the COVID-19 vaccination, the Ad Council used this technology, carefully measuring whether it could increase engagement and vaccination rates. Through a series of live campaigns, starting with just nine messages which grew into hundreds, the team achieved significant gains in engagement by using AI personalization. This research not only saved lives and reduced hospitalizations, but also dramatically demonstrated the power of AI personalization.

After a Stroke, AI Helped Me Learn to Write Again Mukul Pandya, The Wharton School, University of Pennsylvania

Every year, more than fifteen million people around the globe suffer strokes. Those who survive often find themselves in a dark place once they realize that the resulting impairments may be with them for a long time. This mental anguish can be just as debilitating as the more direct effects of the stroke. Mukul Pandya, a lifelong writer and editor, describes how recent developments in artificial intelligence helped him to recover both his writing abilities and his sense of self after a debilitating stroke changed his life overnight.

Section II: Automation
AI + Human Is the Essential Formula for Customer Engagement Chetan Dube, Amelia

In the future, customer engagement in business will rely on a combination of digital and human work. Companies often overlook the necessity of investing equally in both sides of this equation. AI systems and humans must work together as a cohesive team to create personalized user engagement and customer support. Such collaborations can engender repeat business, brand loyalty, and high customer satisfaction over time. Business leaders must therefore be prepared to step out of their comfort zones and make the attempt. It is also essential that they properly prepare their employees for these changes, involving them in the planning, designing, and deployment of AI collaborative projects so they recognize the model's potential to improve not just the business but their own jobs and livelihoods. Companies should trust the technology and their employees in equal measure so they can work together harmoniously, producing impactful customer engagement. Any strategy that doesn't account for the profound contributions of both is out of tune and out of touch.

The Leader's Strategic Mindset: A Key Factor for AI Success Pernille Rydén, IT University of Copenhagen; Torsten Ringberg, Copenhagen Business School; Omar A. El Sawy, Marshall School of Business, University of Southern California

For businesses that aim to succeed with artificial intelligence, their leaders' strategic mindset is a key factor. Businesses are increasingly using AI to transform their processes and create new ways of engaging with customers. An analysis of more than 1,000 leaders' mindsets reveals the danger of a disconnect between the breathless use of the technology and the strategic mindset behind it. Leaders risk falling back on old assumptions about customer engagement and projecting them upon the use of AI. If this mindset doesn't fit, neither the enterprise nor its customers will draw value from the investment in AI. Finding the right 'why' depends on leaders' goals, their intended customers, and their actual AI needs. Our strategic mindset framework, along with examples

of AI/mindset combinations can help managers to find the mindset that best fits their AI customer engagement needs.

Section III: Predictions
White Glove Service: AI in Wealth Management Raises Client Engagement at Scale
Brian Lincoln, IBM Consulting; **Robert Grant,** IBM Consulting; **Suresh Iyengar,** IBM Consulting

The rush of new investors to the markets over the past two years and the growing pool of affluent individuals seeking advice have created a significant strain on wealth manager operations. Clients want to feel that firms know them and expect the same digital convenience they experience in other parts of their lives. However, firms are challenged by the reality of physical documentation requirements and paper intensive processes, with back offices under pressure to automate all while clients seek a seamless digital experience. Leaders have started to apply artificial intelligence (AI) and machine learning (ML) to dramatically enhance intelligent document processing (IDP) to straight through process client requests. The same type of machine learning models used for IDP to transform client experience, reduce processing times, limit or eliminate errors, and provide regular status updates can be used to drive substantial personalization. The generative capabilities of foundation models offer a range of possibilities for wealth advisors to personalize their interactions with clients while driving engagement at scale. Even simple applications of AI to documents can produce impactful results, reducing cycle time, reducing costs, increasing capacity, and improving client engagement.

Resurrecting Jimi Hendrix: The Power of AI to Expand Consumer Engagement Through Musical Fan Cultures
Alan Schulman, UpperRight, **Stacey Lynn Schulman,** HI: Human Insight

Many of history's greatest composers, musicians, and performing artists have created their own unique, instantly recognizable sound. How might their music have evolved had they lived on and continued to create? In recent years, artificial intelligence (AI) and machine learning (ML) have been used to attempt to replicate, reimagine, and expand on what these historic artists created. This study explores how listeners react to AI-generated music that endeavors to interpret and advance the signature sound of several world-renowned artists. Through an analysis of widely available consumer sentiment as well as our own independent research, we explore the delicate boundaries between art and technology and propose a framework for assessing how AI-generated music provides the aesthetic and commercial saliency consumers expect from human musicians. This technology has far-reaching implications for consumer engagement with brands that invest in associations with musical fan cultures and sonic branding.

Wizenoze: The Value of Engaging the Customer's Customer
Stefano Puntoni, The Wharton School, University of Pennsylvania

The Internet offers an incredible wealth of data, but finding relevant, reliable, and readable information for students from primary to graduate school is difficult. Yet, search engines like Google cannot match content to the reading skills of the user and do not always yield results suitable for quality education. The Dutch educational technology startup Wizenoze offers a better way. Founded in 2013, Wizenoze has found a way to curate information on the Internet for children and students, ensuring that what they see is age appropriate, true, safe, and suited to their reading level. After its initial launch, the company's leaders realized that they had considered the needs of educational institutions, but they had not listened to end users: students and teachers. In consultation with these groups, they realized that even though Wizenoze gave students access to a smaller, curated trove of information than the open Internet, many students still found searching for specific information difficult. They wanted an easier way to get the best information about their topic. The company's solution is a demonstration both of the importance of considering end users and of AI's ability to personalize service

at scale in order to increase user engagement and satisfaction.

Developing a Trustworthy AI Rating System and Its Impact on Customer Engagement

Jennifer Shkabatur, Reichman University; **Alex Mintz,** Reichman University

Studies have shown that customers who trust products and services are more engaged, which makes their experience and satisfaction higher. But users are in the dark about the trustworthiness of products that influence their everyday choices, such as recommendation engines, chatbots, mobile apps, digital assistants, and the like. They also have no tools with which to assess them. We propose a method for rating the trustworthiness of AI products and their impact on customer engagement, demonstrating its utility through three real-world applications of AI. Our method allows users to practically test AI products against commonly accepted trust parameters such as human agency and oversight, technical robustness and safety, privacy and data governance, transparency and explainability, non-discrimination and fairness, societal and environmental well-being, and accountability. This method can be used by the business community, investors, regulators, ranking agencies, and customers.

Section IV: Customer Insights

Mars' ACE: Using AI and Behavioral Data in Ad Testing with High Correlating Sales
Laurent Larguinat, Mars, Inc.

Advertising testing is a time-honored tradition used by marketers to gain insights into consumers' responses to ads. Typically, they use standard survey methods in which respondents are asked about ad-recall, ad-preference, or future purchase desire to gauge reactions before the ad is launched. But traditional ad testing is limited by its declarative, survey-style approach; today's marketers need an alternative solution to drive sales and create a sustainable competitive advantage. Mars developed its own pro-

prietary method of assessing effectiveness by using consumer behavioral methods and applying artificial intelligence (AI). We used behavioral data, including eye-tracking, facial reactions, skin conductivity, EEG, and more, applying AI algorithms to analyze the data and ultimately link benchmarks to sales success. The result is 85 percent accuracy in predicting whether ads would lead to a sales.

Machines that Dream: How AI-Human Collaborations in Art Deepen Audience Engagement
Refik Anadol, Artist; **Pelin Kivrak,** Refik Anadol Studios and Tufts University

Human engagement at the intersection of the physical and the virtual world is becoming ever more interactive and multisensory. The Los Angeles-based Refik Anadol Studio experiments and collaborates with artificial intelligence (AI) and the metaverse to create immersive artwork that evokes a range of senses and fosters audience engagement. As the use of generative AI models for artistic purposes becomes more prevalent, it inspires essential discussions about where creativity lies in human-AI collaborations. Artists have a range of perspectives on how to use AI and machine learning (ML) to creatively engage audiences.

Unlocking Deeper Insights into Customer Engagement Through AI-Powered Analysis of Social Media Data
P.K. Kannan, Robert H. Smith School of Business, University of Maryland; **Yi Yang,** Hong Kong University of Science and Technology; **Kunpeng Zhang,** Robert H. Smith School of Business, University of Maryland

Brands use social media channels to engage with their customers by posting content to inform them of new products and services, get feedback, and increase sales through influencer marketing. Users interact and engage with these posts by liking, commenting on, or sharing the content. These interactions between brands and users result in a network of links characterizing a social engagement structure, highlighting the common users between brands. AI

techniques can help firms understand the structure of social engagement and derive market intelligence from this rich data source. We used deep learning auto-encoder techniques to process the network data and identify the competitive structures within brands and product-markets and those that cross boundaries to include different products and brands. Through such analysis, companies can anticipate trends in the market that could affect their brands, noting threats from competing brands and opportunities for co-branding by identifying attractive segments of customers to target and ultimately grow their business.

How Artificial Intelligence Can Keep Classical Music in Business
David Serkin Ludwig, The Juilliard School

Audiences have been engaging with what we call classical music for over a thousand years, but the future of the art form is uncertain today as it competes for attention and relevance. Developments in technology, especially artificial intelligence (AI), are quickly changing the way classical music is produced and disseminated. Rather than displacing creators and performers, AI can be a tool to help musicians write, play, and teach more effectively and engage audiences on a far greater scale.

Infosys: AI Helps Build Customer Engagement to Ace the Tennis Game
Navin Rammohan, Infosys; **Mukul Pandya,** The Wharton School, University of Pennsylvania

Eight years ago, Infosys, a global IT services company headquartered in Bangalore, was looking for a global sport through which it could build customer engagement using AI and digital technology. It hoped to reimagine the sport through data and digital experiences. It also hoped that the sport would have a huge following in its primary markets of North America and Europe. Almost a decade later, the company has used AI and other digital technologies to deepen the engagement that key stakeholders – including fans, coaches, players, and the media – have with the game.

Section V: Omnichannel Engagement
AI for Customer Engagement at Google
Google's Customer Engagement leadership team: **Tim Frank, Aastha Gaur, Abheek Gupta, Doris Neubauer, Ian Suttle, Leo Cheng, Natalie Mason, Qiushuang Zhang, Ravi Narasimhan, Roman Karachinsky, Sandeep Beri, Shashi Upadhyay, Tony Li, Vicky Ge**

We spotted significant opportunities to better apply AI/ML research at all customer engagement touchpoints, creating step change opportunities to provide value to our customers, both directly and through Google's representatives. Always working to put our customers first, we offer a prioritization framework, rooted in case studies on the AI we employed, tracing the impact of our AI/ML applications on customer engagement. We learned a number of valuable lessons from this work.

Singapore's Hospital to Home Program: Raising Patient Engagement Through AI
John Abisheganaden, National Healthcare Group, Singapore; **Kheng Hock Lee,** SingHealth Group; **Lian Leng Low,** SingHealth Group; **Eugene Shum,** SingHealth Group; **Han Leong Goh,** Integrated Health Information Systems; **Christine Gia Lee Ang,** Integrated Health Information Systems; **Andy Wee An Ta,** Integrated Health Information Systems; **Steven M. Miller,** Singapore Management University

Singapore's public healthcare system has designed a Hospital to Home (H2H) program to safely reduce inpatient readmissions of elderly patients within 12 months of an initial stay. H2H focuses especially on patients who suffer from multiple chronic ailments and have complex care needs, helping these patients to safely transition back to their homes and communities. It uses an AI prediction model trained on health record and demographic data to perform the initial patient screening for program enrollment. We summarize our observations on how AI can play such a supporting role behind the scenes, enabling a healthcare program that provides direct engagement with every patient.

Commonwealth Bank: Amplifying Customer Centricity with AI

Steven Randazzo, Warwick Business School and Harvard University; **Jin H. Paik,** Harvard Business School and Altruistic; **Yael Grushka-Cockayne,** Darden School of Business, University of Virginia

Commonwealth Bank leveraged data and artificial intelligence (AI) in its operations to maintain a competitive edge by bridging gaps between retail, call center, and digital services. The transformation was driven by the increased availability and standardization of data, which allowed the bank to create new differentiators, such as customization and enhanced interactions. In addition, as AI continues to influence decision-making, the bank commoditized speed, access, and price, placing greater focus on creating unique experiences that customers cannot find elsewhere. The case study revealed that managers developing retail-based AI programs should: increase the adoption of AI and eliminate its appearance of 'magic' by explaining the models and outputs, train workers affected by the AI and ensure that its further use and implementation are guided by their feedback, and consider how the AI can deliver a range of messages that will resonate with customers, not just sales messages.

Singapore's AI Applications in the Public Sector: Six Examples

Steven M. Miller, Singapore Management University

These six examples of Singapore public sector AI applications illustrate different ways of enhancing public engagement across the domains of transportation, municipal services, employment and retraining support, parks and recreation areas, K-12 education, and border security. For such applications it is essential to strike a careful balance between the benefits of enhanced and automated accessibility, responsiveness, and personalization and the risks of lack of public trust, perceived overreach, lack of understandability, and

excessive complexity. Leaders must also select projects with clearly defined use cases which address people's needs rather than pursuing projects just for the sake of using AI. They should also start with basic tasks that help residents to do essential things more easily, building a foundation for moving forward with more sophisticated AI methods and models that offer new services to residents. Most of these examples were the result of steady, iterative efforts spanning multiple years of patient, persistent, and consistent management, along with parallel efforts to establish policies, governance, and technology platforms and products. The government's approach to building and deploying AI capabilities requires engagement at the national level.

Section VI: The Future
3 Visions of the Future of AI for Customer Engagement: 2027 Scenarios

Jerry Wind, The Wharton School, University of Pennsylvania, **Mukul Pandya,** The Wharton School, University of Pennsylvania, **Margherita Pagani,** SKEMA Business School, **Jerry Dischler,** Google

Artificial intelligence is making ever deeper inroads into every aspect of business, society, and our lives. By 2027, how will AI transform the future of customer and human engagement? *Management and Business Review* and the ANA's Global CMO Growth Council organized a forum with this issue's authors and selected CMOs to explore three possible scenarios for the future - optimistic, pessimistic and most likely. Their results comprise a fascinating combination of opportunities and risks, good and bad. Most importantly, participants outlined the steps we should take today to address these possibilities and prepare for 2027. These scenarios were developed before generative AI took the world by storm with the release of ChatGPT and GPT4 towards the end of 2022. But generative AI does not invalidate the conclusions presented in this special issue and in our scenarios. In fact, it turbocharges the transformative effects of AI on customer engagement and on our lives.

Individuals and Libraries:

Join 830,457 Subscribers and Become an MBR Charter Subscriber to Receive a Price Discount for Life

When you subscribe for 2023 at the low introductory price, you will become a charter subscriber, which will entitle you for life to a discount on the list subscription price. All issues of the 2021 inaugural volume are available free of charge. While other journals charge additional fees for classroom use of articles, MBR will make its articles freely available for use in degree program classes at any school whose library subscribes to MBR. Subscribe online at: https://mbrjournal.com/join-mbr/

	Digital	Subscription Prices
		Four Issues Per Year
		2023 Charter Subscribers
Individual	Digital	$30
Libraries: Colleges and universities	Digital	$400
Libraries: Community colleges, colleges with a budget of less than $1 million, nonprofit organizations, and government departments	Digital	$200
Libraries: K-12 schools	Digital	Complimentary

A Premier Journal for Professional Development

Attractive Prices and Customization Options for Corporate Subscribers When Subscribing for All Their Employees

Each annual volume of the *Management and Business Review* (MBR) will include articles on timely and important issues in business and management. Reading these articles is a highly effective way for your employees to keep up to date on the latest ideas in management from preeminent business scholars and leaders. You can try out MBR in 2023 at a low introductory price. If you do, you will also become *a charter subscriber,* which will entitle you for life to a discount on the list subscription price. Digital copies of all three 2021 issues are complimentary.

	Per Person Subscription Price	
	Four Digital Issues Per Year	
	2021	2023 Charter Subscribers
All managers and professionals	Complimentary	$10
All employees who are not managers or professionals	Complimentary	$2

Customize MBR to Serve Your Constituents and Elevate Your Brand

You can customize MBR by adding special content tailored for your company's mission and the needs of your employees, customers, and suppliers. This customization will strengthen your relationship with your constituents and elevate your company's brand. Companies that choose to customize MBR will be listed on our masthead as Partner Companies. For bulk subscriptions and customization, please contact Publisher and Coeditor in Chief Kalyan Singhal at <Ksinghal@ubalt.edu>.

A Special Offer for Business School Deans

Serve Your Alumni and Other Constituents and Elevate Your Brand

We invite you, the deans of business schools, to offer complimentary digital issues of MBR to your alumni, at only a nominal price per alum. You also have the option of customizing MBR with content specific to your schools. By customizing, you can insert content tailored to your school's mission and the needs of your stakeholders, including business and non-business alums of your university, employees of your industry affiliates, and your faculty and students.

This customization will benefit your stakeholders and elevate your school's brand. MBR is also offering special pricing on custom options. Schools that choose to customize MBR will be listed as Partner Schools on MBR's masthead. Please contact Coeditor in Chief Kalyan Singhal at <Ksinghal@ubalt.edu> to discuss pricing and customization options.

With our special pricing, you will also be able to share complimentary digital copies of your customized MBR with your business school alums and with alums from your university's other disciplines, acknowledging that students from a diversity of disciplines and majors ultimately go into business. Likewise, the managers of your industry affiliates will be able to subscribe to your customized MBR.

We encourage you to order MBR for your library. Unlike some journals, which charge for classroom use, MBR will make its articles freely available for use in degree program classes at any school whose library subscribes to MBR.

Share All Three 2021 Issues of MBR with Everyone You Know and Encourage Them to Share with Everyone They Know

The *Management and Business Review* is meant for everyone in the business community, around the world. To get it into the hands of as many people as possible, we request that you share all three 2021 issues with everyone you know. They are available at the link below.

https://mbrjournal.com/2021-3-free-issues/

You can pass along these issues in several ways:

1. Share each issue or the above link by email with everyone you know who might be interested. You might also want to share it with everyone in your *Send* folder or any other folder. Please remember to invite your recipients to share it with their contacts and so on. Here are two links to download instructions for sending emails from your folder:

   https://mbrjournal.com/wp-content/uploads/2020/09/Instructions-for-creating-a-contact-list-to-send-emails-to-your-contacts-from-Outlook.docx

   https://mbrjournal.com/wp-content/uploads/2020/09/Instructions-to-send-emails-to-all-your-contacts-in-Gmail.docx

   It may be a good idea to put yourself in the "To" field and your contacts in the "Bcc" field. That way, you will avoid sharing everybody's email address with everybody else on the list.

2. Post the issue or the above link publicly on your website or on social media and invite your contacts to do the same, creating a distribution chain.

3. Invite the head of your organization to share the issue or the above link throughout your organization.

   Thank you for helping us to spread the word about MBR and strengthen our global business community.

Join MBR
on Social Media

We invite you to join us on social media for the latest updates on the *Management and Business Review* (MBR). We will use our social media to give you current information on the journal and also as a space in which academic and industry readers can engage with one another. We hope these platforms will become a nexus of rich and rewarding discourse about business and management for the expansive community of business students, scholars, and practitioners.

We therefore encourage you to make use of the social media resources we have created for you.

Facebook: Follow us at @MBRJournal or https://www.facebook.com/MBRJournal.

Twitter: Follow us at @MBRJournal or https://twitter.com/MBRJournal.

LinkedIn: Follow us at https://www.linkedin.com/company/management-and-business-review-mbr-journal.

Please use **#MBRJournal** and tag us **@MBRJournal** on LinkedIn, Facebook, and Twitter. We also invite you to use these platforms to highlight your own MBR papers.

For suggestions and questions related to @MBRJournal social media pages, please contact Professor Subodha Kumar, Executive Editor, Management and Business Review at subodha@temple.edu.

Papers to Appear in Future Issues of the *Management and Business Review*

Koen Pauwels, Oliver Koll.
Why Brands Grow: The Powers of Differentiation and Penetration

Qing Li, Christopher Tang.
Unlocking the Value of Innovative Selling: Information and Options

Jeff S. Johnson.
You Move Me: Understanding and Optimizing Job Rotation

Anatoli Colicev, Ashwin Malshe, Koen Pauwels.
How Brands Can Leverage Their Social Media Marketing

Robert F. Bordley, Jeffrey M. Keisler.
Responding to the Risk of Imposed Changes on the Project Budget and Deadline

Morris Cohen, Shiliang Cui, Sebastian Doetsch, Ricardo Ernst, Arnd Huchzermeier, Panos Kouvelis, Hau Lee, Hirofumi Matsuo, Andy A. Tsay.
Putting Supply Chain Resilience Theory into Practice

Nirmalya Kumar.
Countercyclical Marketing During Recessions

Christian Schulze and Daniel Blaseg.
The Road to Failure is Paved with Good Discounts: How Premature Discounts Hurt Entrepreneurs on Kickstarter

Yen-Ting Lin, David F. Pyke, and Jayashankar M. Swaminathan.
Identifying and Overcoming Barriers to Implementing Blockchain in Supply Chains

Satish Nambisan, Yadong Luo.
Managing Risks in Digital Globalization

Satish Nambisan.
Digital Strategies for Managing Global Operations in the New Normal

Ann Majchrzak, Arvind Malhotra.
Getting More from Crowdsourcing: Solving Wicked Business and Societal Problems

Ashish Sinha, Prakash Bagri, Kiran Pedada, Rajendra Srivastava.
E-Commerce Disrupted: Is Social E-Commerce the Future of Online Shopping?

Sunil Mithas, Rajiv Kohli.
How to Manage Technology Debt and Shadow IT for a Synchronized Digital Strategy

Maude Lavanchy, Amit Joshi, Arnaud Chevallier.
Don't Let the AI Hype Undermine Good Decision-Making

Benjamin Lawrence, Jie J. Zhang.
Infusing an Operational Perspective into Franchise Management

Vijay Govindarajan, Anup Srivastava.
What Is a Modern Tech Company? How Does It Differ from a 20th Century Industrial Giant?

Ali Aslan Gümüsay, Thomas Bohné,Tom Davenport.
AI and the Future of Management Decision-Making

Stanley Frederick W.T. Lim, David F. Pyke.
The 4th Channel: Automatic Home Replenishment and its Implications for Supply Chain Cost and Customer Convenience

Kate Vitasek, Daniel Bumblauskas, Jim Groton, Yu (Jade) Chu.
The Rise of the Standing Neutral: An Easy, Economical, and Effective Way to Prevent Contractual Conflicts

James R. Francis, Alok Baveja, Xin (David) Ding, Ann D. Bagchi, Benjamin Melamed, Diane Hill.
Navigating Future Shifts in Healthcare Service Delivery: Three Insights from Supply Chain Management

Kieren Mayers, Tom Davis, Luk N. Van Wassenhove.
Seven lessons to address resource consumption

Arnoud De Meyer, Peter J. Williamson.
Kickstarting an Ecosystem to Innovate

Vikas Mittal, Jihye Jung.
Strategic Management of Corporate Political Activism

WHARTON — Aresty Institute of Executive Education — UNIVERSITY of PENNSYLVANIA

EXECUTIVE EDUCATION

Whatever the Challenge,
Be *WHARTON READY.*

On Campus and Online Programs for Executives

When the only constant is change, your strongest advantage is Wharton. Be ready for greater success with the proven strategies and in-depth business acumen you can only get from our world-class faculty. Choose from a wide portfolio of results-focused programs delivered in a live online format and on campus.

FIND YOUR PROGRAM ›

execed.wharton.upenn.edu/*SUCCESS*

Advertising in the *Management and Business Review*

This special double issue is almost ad free by design. We do publish ads in our regular issues, though. Please contact Coeditor in Chief Kalyan Singhal at Ksinghal@ubalt.edu to learn more.