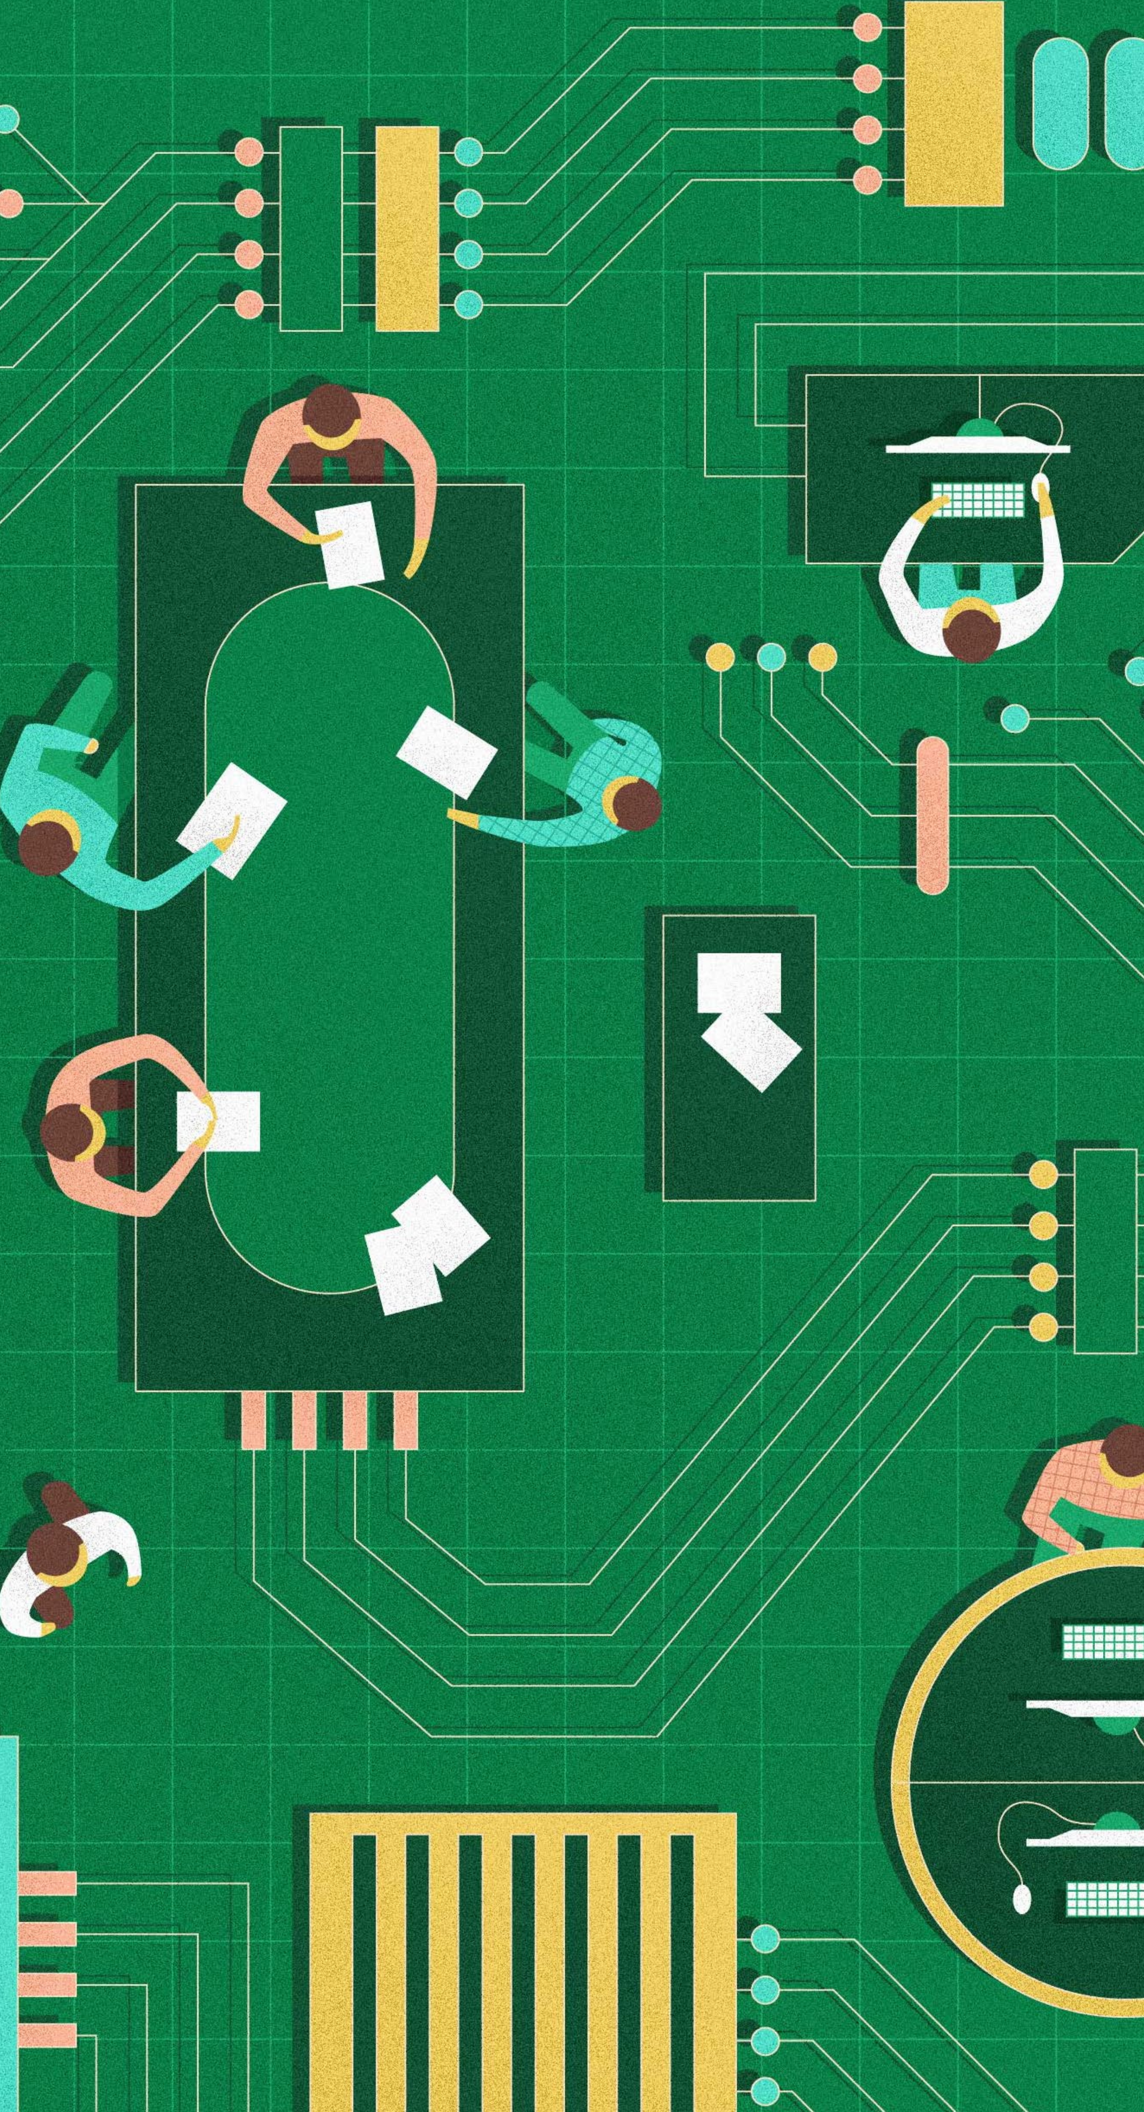

Kellogg Insight Presents:

The Marketing Leader's Guide to Analytics and AI

Kellogg Insight

Northwestern | Kellogg
School of Management

insight.kellogg.northwestern.edu

CONTACT KELLOGG INSIGHT
insight@kellogg.northwestern.edu

SUBSCRIBE TO OUR EMAIL NEWSLETTER
kell.gg/subscribe_2020

KELLOGG SCHOOL OF MANAGEMENT

DEAN
Francesca Cornelli

KELLOGG INSIGHT

FACULTY DIRECTOR
Thomas N. Hubbard

EDITOR IN CHIEF
Jessica Love

SENIOR EDITOR
Emily Stone

BUSINESS AND ART EDITOR
Fred Schmalz

RESEARCH EDITOR
Jake Smith

MARKETING MANAGER
Kevin Bailey

COPY EDITOR
Maja Kos

CONTRIBUTING WRITERS
Susie Allen, Drew Calvert, Glenn Jeffers, Jessica Love, Fred Schmalz, Emily Stone, Dylan Walsh

ILLUSTRATORS
Will Dinsky, Riley Mann, Michael Meier, Michael Morgenstern, Yevgenia Nayberg, Morgan Ramberg, Lisa Röper

KELLOGG EXECUTIVE EDUCATION

Gretchen Baker, Genevieve Burke, Sammi Esterman, Michael Peck, Daniel Underriter

KELLOGG COMMUNICATIONS

Shannon Camperson, Meredith Klein

KELLOGG CREATIVE DIGITAL SERVICES

Kay Meyer, Sofia Park, Marnie Vosper

Hello, marketing leaders.

Data has always played a substantial role in organizations' ability to understand and serve their customers. What's changed in recent years, of course, is the scale and methods with which data can be collected, analyzed, and deployed—all in real time and in response to ever-changing customer needs. If your organization isn't preparing for this new reality, in the words of one of our contributors, it "will be dead sooner than you think."

So it's no wonder that many leaders are feeling anxious: anxious to improve their analytics and artificial-intelligence capabilities, anxious to purchase the right software or hire the right team, anxious to just get started already. But the critical and multifaceted decisions about where and how a company should focus its analytics efforts cannot be outsourced: not to vendors hawking one-size-fits-most solutions and not even to a company's own data scientists. Leaders themselves are ultimately responsible for transforming their organizations. This ebook is for them.

Careful readers will note that missing from this collection is the single answer, the silver bullet, the one right path that will always lead to success. Instead, Kellogg and Northwestern faculty, as well as leading practitioners, share their perspectives on how leaders should think about data, analytics, and AI—and offer practical tips and strategies for confidently using these tools to solve actual business problems.

The first set of articles in this collection addresses broad questions about what, exactly, leaders need to understand about analytics and AI in order to use them effectively. Kellogg's **Florian Zettelmeyer** makes the case for why leaders must develop a working knowledge of data science. **Kris Hammond** from Northwestern's McCormick School of Engineering explains how, even in the absence of a detailed understanding of the technology, executives can work backwards from their business goals to understand how AI can be used in their organizations. Kellogg's **Mohan Sawhney** presents a strategy for determining how an organization should focus its AI efforts at various stages of the customer journey. Kellogg's **Tom O'Toole**, formerly of United Airlines, and **Eric Leininger**, formerly of McDonald's, explore how a company's structure and culture also need to change to accommodate a data-centric focus. And Vanguard's **Jing Wang** shares her own experience building a consumer-analytics team from the ground up.

The second set of articles addresses specific techniques and tools that can help organizations and their leaders up their data game. Northwestern's **Steve Franconeri** shares nine rules for mastering the essential skill of data visualization. Kellogg's **Eric Anderson**, along with Zettelmeyer, explain a counterintuitive strategy for optimizing analytics—one that can be summarized as the need to periodically, but thoughtfully, "mess stuff up" in order to get the data you need to continue learning. Kellogg's **Joel Shapiro** makes the case for analyzing the often-tossed outliers in customer data. Kellogg's **Brett Gordon** shares how experimentation may be the key to better digital advertising. And Kellogg's **Jennifer Cutler** and **Artem Timoshenko** each share a novel technique, backed with research, for using AI to glean customer insights from user-generated content.

Finally, two articles—featuring **Adam Pah** of Kellogg and **Inhi Cho Suh** of IBM—explore how AI is already transforming the field of marketing and examine the questions that organizations are asking, and must continue to ask, about their obligations to their customers and to society.

We are publishing this guide in order to provide leaders with timely and thought-provoking advice for using data to build enduring relationships with customers. But many of these articles have previously appeared, in some form or another, in our publication. To read insights from our faculty as they appear, be sure to subscribe to Kellogg Insight.

We hope you enjoy this collection.

JESSICA LOVE
EDITOR IN CHIEF, KELLOGG INSIGHT

Kellogg Has a Unique Perspective on Analytics

In order to succeed with analytics, businesses need to change how decisions are made. This involves new ways of thinking about analytics, as well as new processes around data collection, experimental design, and planning for the future. Unless businesses commit to making these changes, other pressing business needs will always be prioritized, and years from now, companies will find themselves in the same place, wondering whether their investments in analytics and AI will ever pan out.

You can learn more about Kellogg's perspective in the following pages, but for more in-depth knowledge, consider an executive education course. Learn more here.

Getting Started

Tips & Tools

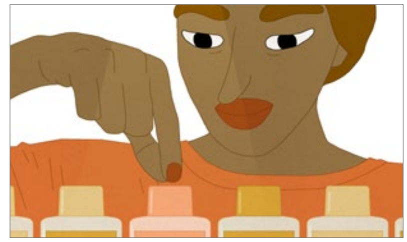

Trends

Getting Started

Why the Most Important Skills in Analytics Aren't Technical Skills—They're Thinking Skills

A working knowledge of data science can help you lead with confidence.

IN RECENT YEARS, DATA SCIENCE has become an essential business tool. With access to incredible amounts of data—thanks to advanced computing and the "Internet of things"—companies are now able to measure every aspect of their operations in granular detail. But many business leaders, overwhelmed by this constant blizzard of metrics, are hesitant to get involved in what they see as a technical process.

Florian Zettelmeyer explains why analytics require managerial judgment.

For Florian Zettelmeyer, a professor of marketing and faculty director of the program on advanced analytics and AI at the Kellogg School, managers should not view analytics as something that falls beyond their purview. "The most important skills in analytics are not technical skills," he says. "They're thinking skills." Managing well with analytics and its close cousin AI does not require a math genius or master of computer science; instead, it requires what Zettelmeyer calls "a working knowledge" of data science. This means being able to separate good data from bad and knowing where precisely analytics and AI can add value.

A working knowledge of data science can help leaders turn data into genuine insight. It can also save them from making decisions based on faulty assumptions. "When analytics goes bad, the number one reason is because data that did not result from an experiment are presented as if they did," Zettelmeyer says. "If you don't understand experiments," he says, "you don't understand analytics."

Start with the Problem

Too often, Zettelmeyer says, managers collect data without knowing how they will use it. “You have to think about the generation of data as a strategic imperative,” he says. In other words, analytics is not a separate business practice; it has to be integrated into the business plan itself. Whatever a company chooses to measure, the results will only be useful if the data collection is done with purpose.

Like all scientific inquiries, analytics needs to start with a question or problem in mind. Whether it is a software company that wants to improve its advertising campaign or a fast-food company that wants to streamline its global operations, the data collection has to match the specific business problem at hand. “You can’t just hope that the data that get incidentally created in the course of business are the kind of data that are going to lead to breakthroughs,” Zettelmeyer says. “While it is obvious that some kinds of data should be collected, customer interactions have to be designed with analytics in mind to ensure that you have the measures you need.”

Analytics vs. Artificial Intelligence

It turns out there’s not a single, agreed-upon way in which these terms are used. Ask a dozen data scientists, and you’ll get a dozen different answers. (Ditto when it comes to definitions for “big data” or even “data scientist.”) But here’s how Zettelmeyer sees it.

Analytics: The science or process of transforming data into knowledge. This is how we derive useful conclusions from data. Analytics can include processes such as machine learning, but it can also include conducting experiments or running statistical analyses that allow you to make causal inferences.

Artificial Intelligence: When automation allows systems to perform “intelligent” actions such as learning, creating knowledge, making inferences and decisions, and solving problems. In Zettelmeyer’s view, AI is so exciting because it scales analytics.

Nor can managers rely on data scientists to take the lead. Ultimately, it is the manager's job to choose which problems need to be solved and how the company should incorporate analytics into its operations. Executives, after all, are the ones who have to make decisions; therefore, they should play a central role in determining what to measure and what the numbers mean to the company's overall strategy.

Understand the Data-Generation Process

"There is a view out there that because analytics is based on data science, it somehow represents disembodied truth," Zettelmeyer says. "Regrettably, that is just wrong."

So how can leaders learn to distinguish between good and bad analytics or AI? "It all starts with understanding the data-generation process," Zettelmeyer says. "You cannot judge the quality of the analytics if you don't have a very clear idea of where the data came from."

Zettelmeyer says most managers share a common behavioral bias: when results

Analytics vs. Artificial Intelligence

Let's say an advertiser is bidding to show an online ad to a customer. They might use mass-scale experimentation to determine how effective an ad will be for certain types of customers and then create a model that automatically scores every potential target in terms of their likelihood of responding in order to calculate a bid. The model itself would be developed and updated offline. But the actual scoring and bidding could be done in real time.

"It's this crazy combination of using mass-scale experimentation, using very sophisticated machine learning, and using automation in order to create something that, if performed by humans, would be considered intelligent," says Zettelmeyer. "The future of learning in marketing is going to come from some form of marrying these together."

are presented as having been achieved through complicated analysis, they tend to defer to the experts. “There is a real danger in managers assuming that the analysis was done in a reasonable way. I think this makes it incredibly important for managers to have a sixth sense for what they can actually learn from data.” To make informed decisions, he says, it helps to take a step back and establish some fundamentals.

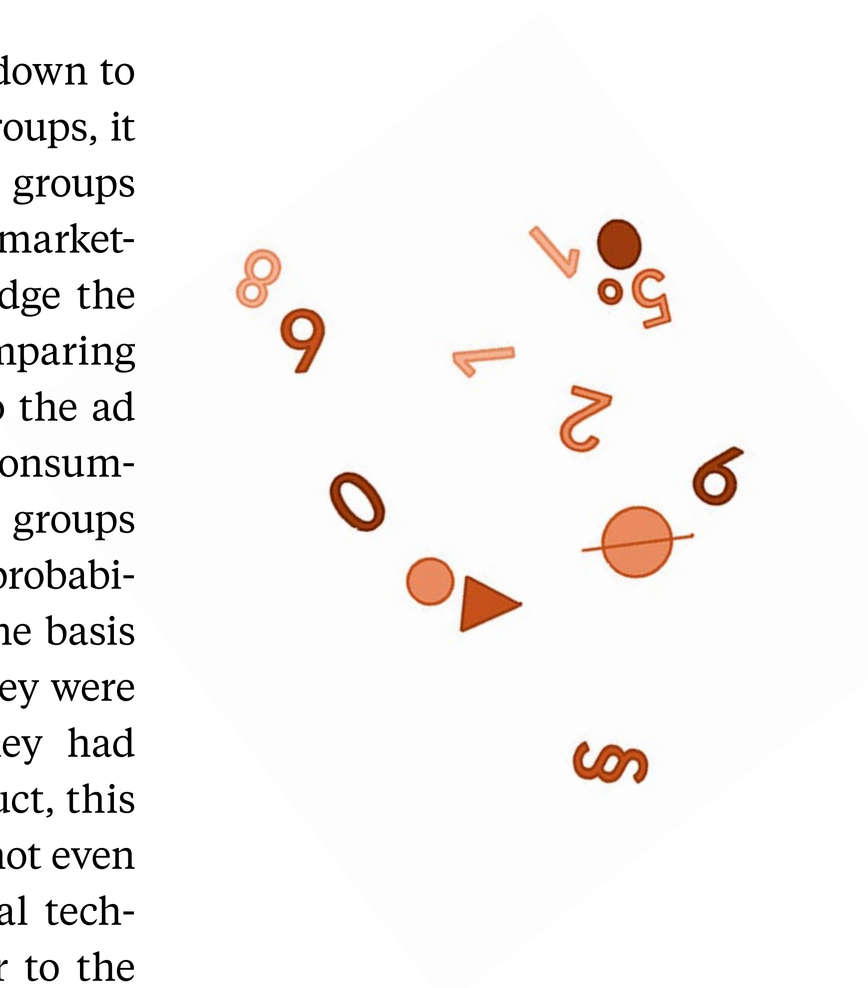

Because an analysis often boils down to making comparisons between groups, it is important to know how those groups are selected. For example, a marketing department may want to judge the effectiveness of an ad by comparing consumers who were exposed to the ad with those who were not. If the consumers were selected randomly, the groups are what data scientists call “probabilistically equivalent,” which is the basis for good analytics. But if, say, they were exposed to the ad because they had shown prior interest in the product, this will lead to bad analytics, since not even the most sophisticated analytical techniques could provide an answer to the basic question: Was the ad truly effective or was the consumer already interested?

“Algorithms are limited by the data you feed them,” says Zettelmeyer. “If opportunistic data created in the normal course of business are polluted, what you can learn is very limited.”

This is not just a marketing problem. Take, for example, a hospital that wants to replace its ultrasound machines. Thanks to advanced wireless sensors, the hospital is able to measure in the course of business exactly how long it takes to perform an exam using the new devices, a metric that would help it decide whether to switch over for good. But the data show a surprising result: the new device is taking longer to use than the older one. What the hospital had not accounted for was a preexisting difference between two groups of technicians: novice technicians and experienced technicians. It turns out that more novice technicians, who were naturally slower than the experienced ones, were choosing to use the newer device, and this skewed the data.

“The problem,” Zettelmeyer says, “is one of confounding technician experience with the speed of the device.” Again, analytics failed because it overlooked fundamental questions: What makes technicians choose one machine over the other? Is everything about the usage of the two machines comparable? And if not, was the correct analytics used to correct for that?

Understanding the data-generation process can also uncover the problem of reverse causality. Here, Zettelmeyer points to the case of a company deciding whether or not to limit promotional emails. The data reveal that promotional emails are extremely effective: the more emails a customer receives, the more purchases they are likely to make. But what is not apparent in the data is that the company is following a piece of marketing wisdom that *Reader's Digest* hit upon decades ago: loyal customers—people who bought more recently, buy more frequently, and spend more on purchases—are more likely to buy again when they are targeted. So rather than the number of emails driving the amount of sales, the causality actually works the other way: the more purchases customers make, the more emails they receive. Which means that the data are effectively useless for determining whether email drives revenue.

Use Domain Knowledge

In addition to making sure that data are generated with analytics in mind, managers should use their knowledge of the business to account for strange results. Zettelmeyer recommends asking the question: "Knowing what you know about your business, is there a plausible explanation for that result?" Analytics, after all, is not simply a matter of crunching numbers in a vacuum. Data scientists do not have all the domain expertise managers have, and analytics is no substitute for understanding the business.

Consider an auto dealership that runs a promotion in February. Based on a rise in sales for that month, the dealer assumes the promotion worked. "But," Zettelmeyer says, "let's say what they were trying to sell is a Subaru station wagon with four-wheel drive, and they completely ignored the fact that there was a giant blizzard in February, which caused more people to buy station wagons with four-wheel drive." In cases like these, he says, having the data is not enough.

This step becomes even more critical as an organization's analytics capabilities increase, and analyses become more powerful—and opaque. Algorithms that rely on neural networks, for instance, are powerful enough to detect complicated statistical relationships among variables, but because they are "black boxes," the nature of those relationships isn't easy to conceptualize. This can increase the tendency to blindly trust whatever results the algorithm spits out, rather than asking common-sense questions about whether wonky results are a function of how the model was built or trained.

"If you buy into this idea that AI ultimately scales analytics, please don't think that it means you're off the hook when it comes to understanding how to create knowledge and what you can learn from data," says Zettelmeyer.

Know It—Do Not Just Think It

As Zettelmeyer sees it, decision-making in the business world is being revolutionized in the same way that healthcare is with the widespread adoption of "evidence-based medicine." As advanced analytics and AI bring about this revolution, managers with a working knowledge of data science will have an edge. Beyond being the gatekeepers of their own analytics, leaders should ensure that this knowledge is shared across their organization—a disciplined, data-literate company is one that is likely to learn fast and add more value across the board.

"If we want data analytics and AI to succeed, everyone needs to feel that they have a right to question established wisdom," Zettelmeyer says. "There has to be a culture where you can't get away with 'thinking' as opposed to 'knowing.'"

Developing such a culture is a big challenge for leaders. Organizations are rarely willing to admit the need for change, and few managers feel confident enough to lead with analytics. This, he says, will have to change.

"Can you imagine a CFO going to the CEO and saying, 'I don't really know how to read a balance sheet, but I have someone on my team who is really good at it.' We would laugh that person out of the room," Zettelmeyer says. "And yet I know a whole bunch of people in other disciplines, for example, marketing, who, without blinking an eye, would go to the CEO and say, 'This analytics stuff is complicated. I don't have a full grasp on it. But I have assembled a crackerjack analytics team that is going to push us to the next level.' I think this is an answer that is no longer acceptable."

> There has to be a culture where you can't get away with 'thinking' as opposed to 'knowing.'
>
> **FLORIAN ZETTELMEYER**

Featured Faculty

Florian Zettelmeyer, Nancy L. Ertle Pofessor of Marketing at Kellogg

Learn more from Florian Zettelmeyer in our **Executive Education programs**.

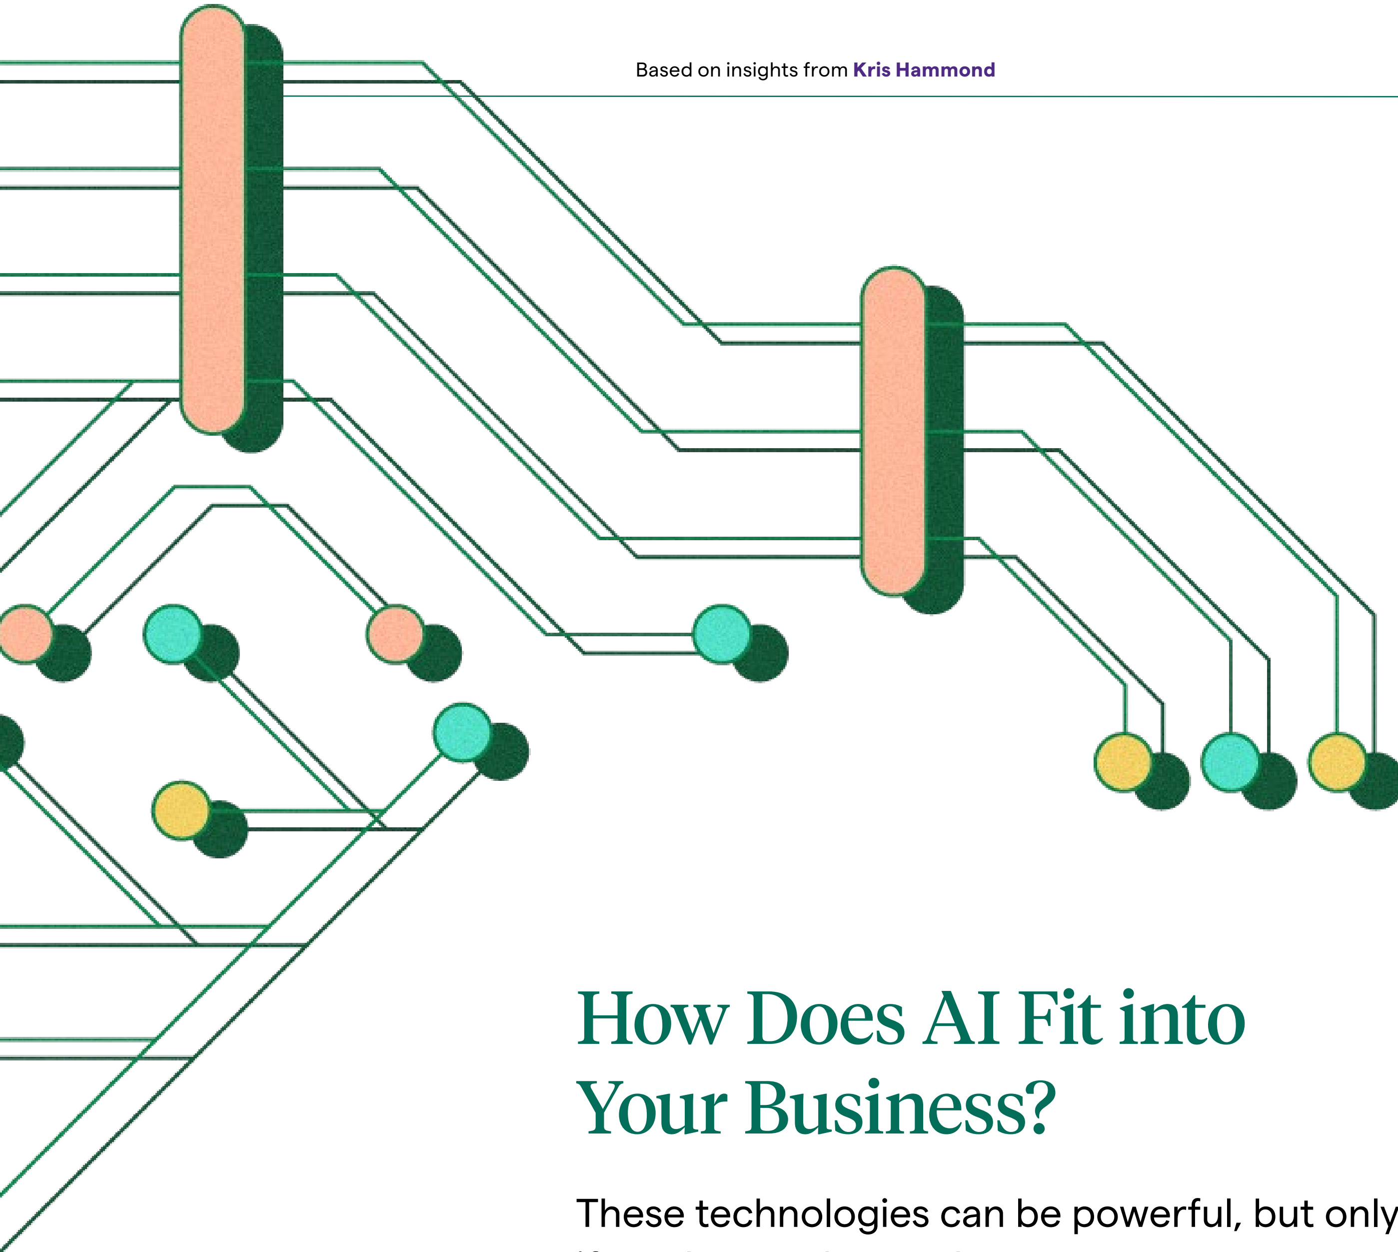

How Does AI Fit into Your Business?

These technologies can be powerful, but only if you know what problems you want to solve.

THE FIRST THING KRIS HAMMOND WANTS business leaders to know about AI is that AI isn't a "thing." Not a single one, anyway.

"It's not a single algorithm, it doesn't use one kind of data, it's not even a single functionality," he says. "Machine intelligence covers a spectrum of functionality, in much the same way as human intelligence covers a spectrum of functionality."

Hammond is a professor of computer science at Northwestern's McCormick School of Engineering. He is also cofounder and scientific advisor for Narrative Science, which uses an AI-powered tool to turn data into stories that can be easily read by real people. And he speaks regularly with executives around the country—including at Kellogg—about artificial intelligence.

In his experience, many executives, while eager to explore the impact AI could have on their business, have a fundamental misunderstanding about just what it is. They view AI as a single magic bullet that can make a firm more efficient or profitable or resistant to disruption. They don't yet see it as a collection of technologies, operating on a wide range of data types, and designed to address a huge swath of problems, from recognizing what is happening now, to predicting what will happen next, to aiding in decision-making.

As such, they are asking a lot of questions about AI—but they tend to be the wrong questions.

"Execs will often ask, 'What should our AI strategy be?' That's just like asking, 'What should our shovel strategy be?'" says Hammond. "But it really depends. How many holes do you need to dig? How deep do those holes need to be? How quickly do you need to dig those holes? Do you dig those holes locally? Regionally? Globally? Are there rules and regulations around digging holes? Do you have performance metrics with regards to digging holes?"

Instead of seeking out a single perfect tool to harness, or even a single right path toward technological success, Hammond argues, these leaders need to do what they do best: think carefully about the business

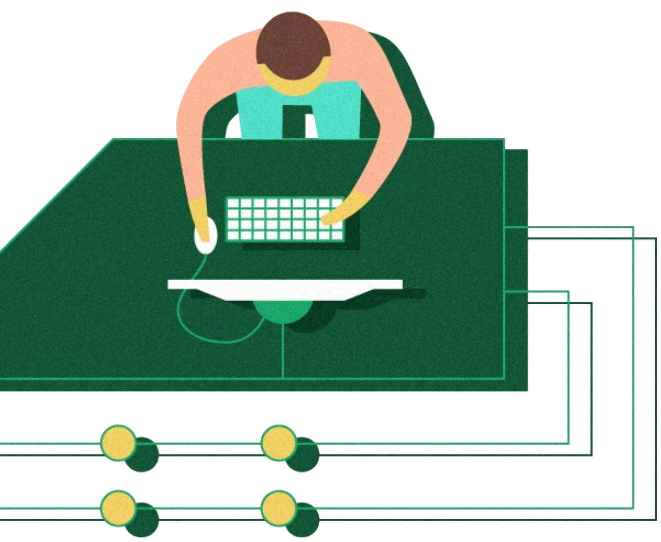

problems that they need to solve, and let these problems drive any decisions about how to move forward.

“You can ask yourself, ‘What are the functionalities that I actually need? Do I need to look at text and understand what people have said? Do I want to predict what people are going to buy? Do I want to understand, given what people are buying, what features actually participated in that?’” says Hammond. “Until you do that, you’ll end up flailing about in these open and empty problem spaces as opposed to solving specific problems for specific reasons with specific pieces of technology.”

Developing a Functional Knowledge of AI

To some executives, following this advice may seem like entangling oneself in a hopeless chicken-and-egg problem. Sure, they need to examine their business in order to find problems that AI might solve—but don’t they also need to learn enough about AI to understand the possible business problems that AI is uniquely equipped to handle?

Perhaps. But Hammond argues that the level of knowledge required to make smart decisions about AI is different than one might imagine. He offers up cellular technology as a metaphor. Most of us have never heard of cell-signal encoding; fewer could pass a test on how antennas work. Yet we're more than comfortable using our cell phones: we know all the things they can do for us, as well as the practical limitations on their utility, such as the fact that they will die unless we occasionally charge them.

What does this look like in the context of AI? Well, executives *do not* need to understand the algorithms that underlie gradient descent or backpropogation. But they might benefit from what Hammond calls a *functional* understanding of some of the most popular AI technologies. So they should understand that a common technique called "deep learning" is extraordinarily powerful when it comes to recognizing patterns, for instance, but that it has some important constraints as well. Notably, its algorithms require a *lot* of data to train—hundreds of thousands of examples. And deep learning is opaque: once it has learned to recognize patterns, it cannot tell you what features it is picking up on or emphasizing as it recognizes and categorizes. That means that, from a business perspective, it shouldn't be used in those places where the ability to explain or audit decisions is required.

It's this level of information that should influence decisions about whether a piece of technology might be appropriate for a given task.

"So you look at that and ask, 'What does that mean, from a functional point of view, in terms of my business?'" says Hammond. "Well, if you're making decisions about whether or not to give somebody a mortgage, you might be able to use historical data and train up the deep-learning system, and that will give you a yay or a nay. But when you ask that system, 'Why did you just deny someone a mortgage?' the system can't tell you. You've got to ask yourself if that is okay."

Questions about whether a particular technology is appropriate for a problem aren't hypothetical; they can have real—and really powerful—implications for organizations, their customers, and their employees. For example, Amazon used deep learning to train a resume-filtering system on data that included the original resumes of staff members, as well as their subsequent performance reviews. The company's hope was that the filtering system would pick up on signals—particular skills or experiences—that would indicate future success at the company.

But when they started to feed resumes through their system, they noticed a problem.

"It turned out the system was gender biased," says Hammond.

Learn more about how bias creeps into AI when businesses aren't looking.

The system had picked up on signals Amazon never intended it to learn from, and it learned to penalize female candidates. But, Hammond stresses, it wasn't the actual algorithm that was biased—it was Amazon itself. Because the company used data from performance reviews to train the system, and because women similarly qualified as men were sometimes reviewed more poorly, the system logically learned to disfavor the resumes of female applicants.

"I love this example because it shows how machine learning and scale expose things," says Hammond.

It also shows why having a functional understanding about various AI technologies can be so critical. Knowing that deep learning has this potential to amplify any biases reflected in the data can help companies decide when and how to use it, as well as design safeguards to mitigate any negative consequences.

"The kind of knowledge you need, it's not deep," says Hammond. "It really comes down to, for any given technology: What does it do, what are its requirements, what can you expect of it, and where are the places where you can make mistakes?"

The Power of Knowing Your Own Constraints

Still, just because the level of knowledge required isn't deep or particularly technical, it isn't easy for leaders to come by.

They won't get it directly from the companies selling AI-powered tools. Individual vendors have no incentive to provide potential customers with a fundamental understanding of what it is they do—especially now, with competition in the space proliferating. From the vendors' perspective, it is far better to let customers believe their tool works almost as if by magic. Bigger tech companies have even less incentive, as they are likely trying to sell you on not just one tool but a suite of tools: not just Microsoft Office, but also Power BI and Azure, and so on.

"Who is incented to explain things in a way that allows companies to make genuinely deep, informed decisions? It turns out nobody is, because it often will get in the way of their goals," says Hammond. "It's easier if you find yourself in a relatively comfortable position where you know how things work and nobody else does."

This leaves executives—particularly those without ready access to data scientists—in an admittedly tough spot. But instead of going straight to vendors and expecting them to provide a candid analysis of the pros and cons of their tool and the technology that underlies it, Hammond advocates working backward from the business problems at hand.

"You have to pull back and ask, what are the things that our business needs to accomplish? What are our business goals?" says Hammond.

And then, within those goals, what are the nuanced issues that your company is likely to face? If you want to build a system that can predict which individuals to hire or whether a potential client will pay back a loan, what do you want to base that system on? Do you want to mirror the rules your firm currently uses, or do you want to take a different approach? Are there historical data that are relevant, and what limitations might those data have? Do you want your system to interact directly with clients, or with your employees instead?

"You can just keep asking questions, and you can build yourself up a nice set of constraints that have nothing to do with technology, but everything to do with what you want the business to become," says Hammond. "Now when you talk to a vendor, you have a set of questions you can ask. And when that vendor says, 'Oh no, this is too complex. You're never going to be able to understand it,' you can say, 'Yeah, you need to be able to answer these questions.'"

This general approach is fairly straightforward to adopt—but it does require something that many executives, particularly those in the highest echelons of the organization, lack: the confidence to ask basic question after basic question, regardless of how naïve or ignorant it may make them feel.

> "People feel like it's an Emperor's New Clothes moment...but they should not feel embarrassed about asking questions."
>
> **KRIS HAMMOND**

"People feel like it's an Emperor's New Clothes moment," says Hammond. "Like, 'if I ask the question, maybe it's showing them that I'm not smart enough to have understood things.' But they should not feel embarrassed about asking questions like this."

The Product Definition Comes First

Leaders with technical expertise at their disposal are in an even better position to effectively use AI to solve problems. In this case, Hammond advises that executives first go through the same exercise of determining the constraints around the technology that they need, but then develop a partnership with someone who can act as a technology strategist.

“That partnership combines the functional level of ‘here’s what we need’ with a technical level of ‘here’s what can be done,’” says Hammond.

With both perspectives accounted for, leaders have the ability to build out the specifications for what it is they want, which Hammond calls the “core product definition.” This means characterizing their goals at the level of the business and defining core products at the level of, for instance, knowing they will need to predict machine downtime on the factory floor, qualify customers before they spend time trying to sell to them, or recognize faulty products as they come off the line. Only once these definitions are in place should the team broaden to incorporate expertise in AI development, UI/UX, data science, and software development. “The product definition comes before anyone starts thinking about real implementation,” says Hammond.

This approach has some helpful built-in guardrails, he explains. “You never go for just the shiny thing. You always track what is valuable to the business. And you’re always aligned with what’s doable from the point of view of the data that you have and the data you can develop.”

What leaders should *not* do, however, is delegate decisions about technology to a single person whose performance is evaluated on technology adoption rather than the firm’s performance. “Because then what they’ll do is they’ll look for places to use the technology,” Hammond says, “as opposed to really focusing on solving the pressing problems of the moment, and pushing the company forward.”

Above all, leaders should stop expecting magic bullets and be skeptical of anyone who promises anything before first answering questions about their tools in simple language.

"Most people in business are living in a world where they've been told there are miracles," says Hammond. "There aren't, but that's okay. There's great tech out there. Just keep taking your business seriously, and things will flow from that."

Featured Faculty

Kris Hammond, Bill and Cathy Osborn Professor of Computer Science at Northwestern University's McCormick School

Related Resources

Watch: Humanizing the Machine with Language (TedX Chicago)

Watch: Machine Consciousness (TedX Northwestern)

How to Use Automation and AI to Radically Transform the Customer Experience

Is your company prepared to deliver the millions of personalized interactions that your customers expect?

IT'S NO SECRET THAT CUSTOMERS EXPECT better service from companies these days.

Over the past decade, tech-savvy firms like Amazon, Uber, and Netflix have conditioned customers to expect service that is personalized to them, accessible on demand across multiple channels, and responsive to their needs at that exact moment.

Now, no matter what you sell, you're on the hook to deliver this kind of experience. "Your customer's expectations are not being set by your industry. They are being set by the best of the best of the best," says Mohan Sawhney, a clinical professor of marketing at Kellogg.

So, across a range of industries, companies are responding. For the best of them, this is not just an effort to stay relevant to their customers, but to find new business opportunities as well. These companies are adopting a mentality that Mohan Sawhney describes as "care is the new commerce."

"What does 'care is the new commerce' mean? It is about eliminating artificial silos between marketing, sales, and customer service," says Sawhney. "After all, a customer is a customer. They are not 'marketing customers' or 'service customers.'"

"Let me give you a scenario shared with me by the former chief digital officer at Nike," Sawhney continues. "Nike has collected 290 million profiles of customers. One of those customers tweets to Nike saying, 'Hey, I'm running the Boston marathon on Monday. It's Thursday and I'm in Florida, and I'm traveling to Boston, but my shoes haven't arrived. I'm getting nervous.' The representative at Nike responds, 'I'm so sorry. We'll expedite the order, but we notice that you're living in Florida. And it's projected to be 29 degrees when you run on Monday. Are you prepared with the right clothes?' The customer says, 'Well, I had no idea.' The representative goes on to say, 'Here are some layers and thermals that you might be interested in. We recommend a layering approach. Let me show you samples of what I think will work for you.' The customer picks out a few, and the rep says, 'We'll put all this together. Tell us the hotel you're staying at in Boston, and it will all be there for you in time for the marathon.'"

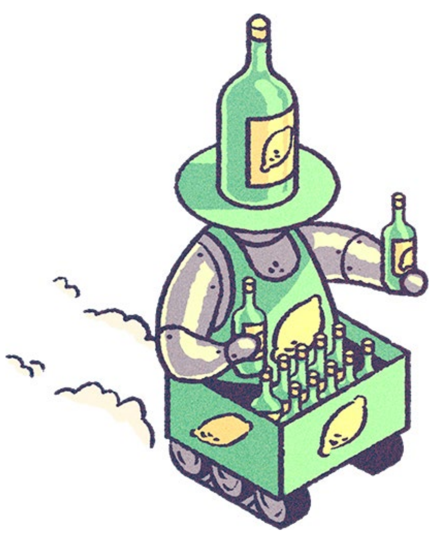

So what began as an interaction with a customer who called in to complain ultimately resulted in a $500 sale.

But this seemingly serendipitous encounter was not, of course, serendipitous at all. The interaction only occurred because Nike's customer service representative had access to relevant information about the customer and was able to look holistically at the customer's needs at that moment and seamlessly switch from being a customer care representative to a sales representative.

It occurred, in other words, because the organization had adopted the "care is the new commerce" mindset.

He adds, "In the future, by the way, this whole conversation would take place with a chatbot. You don't even need a human being."

But organizations that want to transition to this new approach need to do much more than promote it as a mantra. For many, it will take a massive undertaking in three areas: creating a customer-data platform to get all customer data in one place; building a unified system of engagement that allows customers to interact across all digital and physical channels; and leveraging analytics and artificial intelligence to personalize the customer engagement.

Sawhney shares some of the ways companies need to be approaching these challenges in order to deliver the thousands, or millions, of nimble, bespoke customer interactions that will drive sales.

Know Where to Focus Your Investments

In order to provide personalization and responsiveness at any kind of scale, organizations will ultimately need to lean heavily on automation and artificial intelligence. But these tools are not one-size-fits-all, and they must be used thoughtfully to address actual needs at important moments in the customer experience.

This means companies need to determine where to focus their investments. As a good first step, Sawhney recommends that companies map out what he calls the Customer Experience DNA.

Most companies think about the customer experience as a sales funnel: the process of converting prospects into customers. However, they also need to visualize the customer experience from another perspective: the steps that the customer goes through in their journey as they make a purchase decision.

To Sawhney, these two processes—the customer journey and the customer-funnel-management journey—can be visualized as two interwoven strands. He proposes the metaphor of the DNA, with its double-helix structure, as the basis of a framework called the Customer Experience

DNA (CxDNA). The first strand in the CxDNA represents the customer's side of the interaction, while the other strand represents the organization's side.

The Customer Experience DNA (CxDNA)

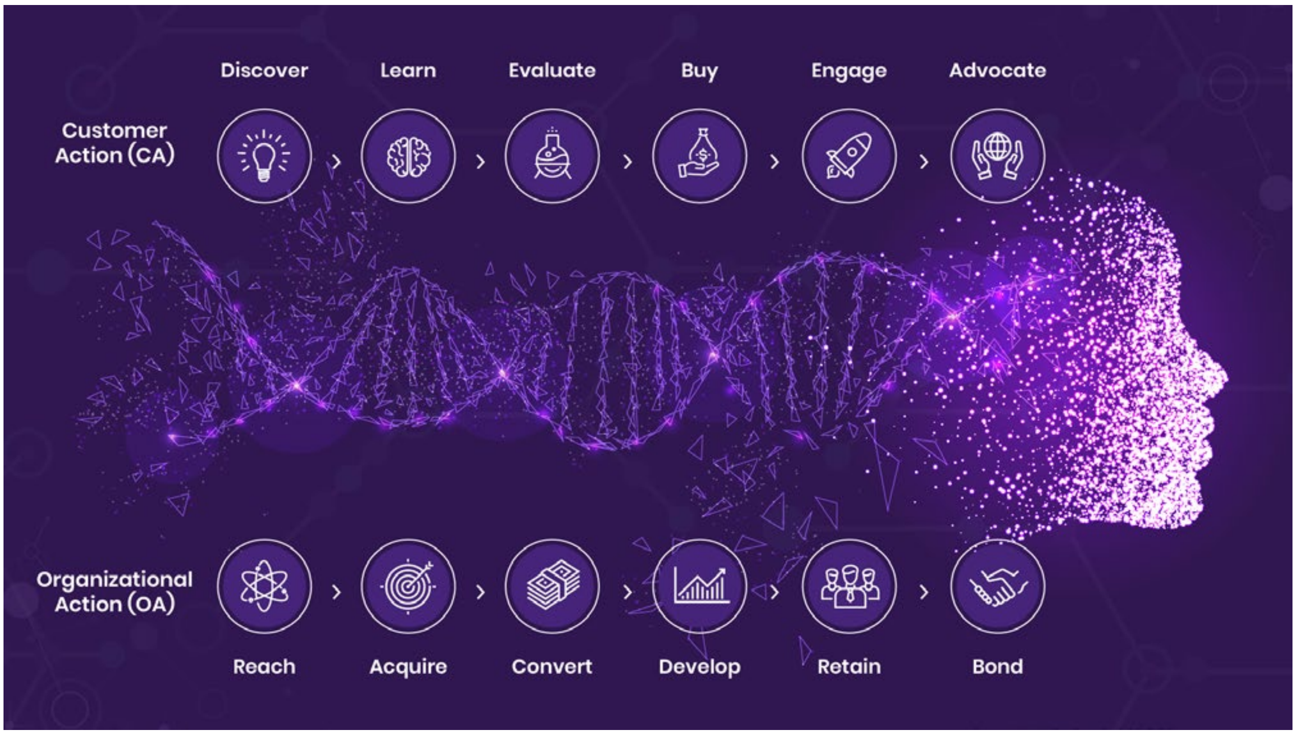

Sawhney advises that companies go through the exercise of exhaustively detailing—both qualitatively and quantitatively—what each stage of the customer experience looks like from the customer's perspective. This will help to clarify what information and content will be most useful for customers at each stage. This analysis in turn informs the organizational actions needed to nurture customers across the customer journey.

"This becomes an organizing framework," he says. "Because once you know what stages customers are going through, you know what to do in order to facilitate their experience. The CxDNA also sets up the use cases for applying artificial intelligence to personalize the information and content to engage with customers at each step in the journey."

Building out a comprehensive, unified model of the customer experience sounds straightforward enough, but Sawhney explains that it can be surprisingly complicated in practice. Many executives lament the lack of a common database about their customers, for instance. "The frustration that they voice is, 'We don't have the data in one place. We have departments and business units and geographies and silos. And unless we can stitch all the stuff together—and by the way, we don't even have the same definition of a customer across the business units—we don't have a seamless customer ID that we can track through all of those silos.'"

This insight often leads companies to create a single platform to house data across all channels, stages, and interactions—and which can also house an automation engine. These are steps Sawhney strongly recommends that companies take before they begin to think seriously about making any other investments into analytics.

"In order to be able to get full value from the analytics and AI efforts, you do need to put the foundations in place," says Sawhney.

Building a Better Funnel

Once you've built out a unified model of the customer experience and developed a customer-data platform capable of tracking individual customer interactions, you are in a good position to test how to bring "care is the new commerce" to life.

Specifically, this means determining which moments in the customer experience could be personalized, rethought, or otherwise improved, and then rigorously testing which interventions work.

"The decision becomes where along the customer journey to run the company's analytics models. Think of it like: you first need to build a road, and then you need to decide whether to drive on the first mile of the road, or the last mile of the road, or some mile in between," says Sawhney.

As a starting point, he suggests looking for leaks in your customer funnel, as these are the places where you are clearly letting your customers down, and a new approach would benefit both of you.

"Look for areas where you are failing, areas where you've got problems. And that becomes the point of attack for predictive models or analytics applications."

Say you're an app developer. Your marketing campaign generated 100,000 leads—and most of those have signed up for a free version of your product. But very few are shelling out for the paid

Looking to cut through the noise and really reach customers?

Done right, digital-media storytelling has incredible power to create conversations between brands and consumers. Sawhney explains why "transmedia storytelling" is the next frontier in branding.

Read more on Kellogg Insight.

version of your product. Your problem, then, is converting them to paying customers.

"That's the point I would attack. I would build a model to predict the conversion and look at variables that drive conversion and what corrective actions I can take," says Sawhney. This might mean, say, developing custom messaging or promotional offers to encourage users to transition to a paid plan—and then testing to see if they work. Or it might mean gauging users' interest in a new feature in the paid version.

Or, in a different scenario, say you are an established software company with deep market penetration—90 percent of the market uses your flagship software through subscription. But many of those customers are using it sparingly, or for limited jobs. Those customers may feel that they aren't getting the full value from the software and may not re-subscribe. This presents you with a customer-retention problem, which you might address by, for instance, developing personalized recommendations for trying out new features or more fully utilizing the software so that your customers can experience more value from it.

A Work in Progress

Fully transitioning to a "care is the new commerce" mentality will not come easily to most organizations, but that is no reason to delay getting started.

"These models and tools and processes are actually in production today. This is not science fiction. This is today," says Sawhney. "So if you are doing marketing and customer engagement the old way—manually and with siloed systems—and assuming that marketing is all about creativity and intuition and running Super Bowl ads, you will be dead sooner than you think."

But to marketing leaders feeling trepidation about how best to deploy technology, Sawhney advises that there is no single best place to focus your efforts. Instead, it is better to improve the customer experience at the places where you have the biggest problems, gradually working towards a more comprehensive end-to-end view.

It is important to think big, but to start small and to scale fast. "Machine learning and artificial intelligence should be applied not to do one gigantic thing," he says, "but to do hundreds of little things that add up to significant business value."

"Machine learning and artificial intelligence should be applied not to do one gigantic thing, but to do hundreds of little things that add up to significant business value."

MOHAN SAWHNEY

Featured Faculty

Mohan Sawhney, McCormick Foundation Chair of Technology and a clinical professor of marketing at Kellogg

Learn more from Mohan Sawhney in our **Executive Education programs**.

Based on insights from **Tom O'Toole** and **Eric Leininger**

Remaking Marketing Organizations for a Data-Driven World

United Airlines developed a robust analytics program to avoid becoming obsolete. Your organization can do the same.

WE HAVE BEEN HEARING FOR YEARS about how the increasing availability of data is changing how marketers engage with customers. But what does that look like in practice? Are customers seeing improved service, better options, and more responsive offers? And how are these changes hitting companies' bottom lines?

For more than six years as senior vice president and chief marketing officer for United Airlines and president of United MileagePlus, Tom O'Toole has had a finger on the pulse of a radically changing marketing landscape.

O'Toole, an executive director for the Program for Data Analytics at Kellogg and a clinical professor of marketing, sat down with Eric Leininger, a clinical professor of executive education at Kellogg, to discuss what marketing looks like in his industry today—and what it may look like tomorrow.

This interview has been edited for length and clarity.

ERIC LEININGER: We've seen marketing shift to become more personalized, engaging customers in a different kind of relationship than in the past. What's surprising to you about the way marketing is changing?

TOM O'TOOLE: It may not be fully appreciated how advanced data-driven targeted marketing has become.

At United, for example, we are currently

Create a Data-Friendly Culture

Tip 1. Encourage Intellectual Curiosity throughout the Organization

Questions about using data to improve business should be welcomed from all corners of an organization. O'Toole was recently discussing analytics with the senior leadership team of a financial-services company, when its general counsel wondered aloud how predictive analytics could be used to identify likely cases of a certain type of regulatory-compliance problem.

"People didn't expect the head lawyer to be the one to ask how to use predictive analytics to address an issue," O'Toole says, "but he did. And, predicting and avoiding the compliance problems that he had in mind would have real business value in avoiding the costs related to lawsuits, customer settlements, and terminations."

doing a targeted-marketing program, particularly to grow our market share in what are called "key spoke" cities, meaning important cities that are not our hubs. This is the largest such program we've undertaken, targeted to several million individuals. Each individual received, and is now tracking toward, a multi-tiered personalized offer based on her or his actual flights, spending on United, and other factors in the last year.

Even a few years ago, this would have been a much more conventional frequent-flyer program promotion based on customer segmentation and promotional-offer cells. Now it is extremely individualized and extremely targeted—and extremely measurable: we develop personalized offers based on individual customer value and other criteria, deliver those offers to specific individuals through email, digital channels, and social media, and then have the feedback loop to business results at the individual level: Did they book? Did their number of flights with us increase? Did their spend with us increase? And, by extension, did our share of the customer's total flights and flight spend increase?

Create a Data-Friendly Culture

Tip 2. Hire Curious People

Building a culture of intellectual curiosity begins in the hiring process. O'Toole says that conversations about a wide range of subjects can offer useful insights into how a job candidate thinks in terms of data.

"For example, if you learn that someone is interested in what's driving crime rates in a particular geographic area—and that they used data analysis to reveal, understand, and visualize the patterns—well, that may not directly relate to your business, but it can show how their mind uses data to solve problems."

This example isn't hypothetical. We optimize marketing at the individual level on a very large scale. Sometimes I hear people talk about how companies are *going to* be doing this. I think, no, we're *doing* it. And others are doing it as well.

> "Sometimes I hear people talk about how companies are *going to* be doing this. I think, no, we're *doing* it. And others are doing it as well."
>
> **TOM O'TOOLE**

LEININGER: Does a traditional segmentation study, in which you identify six or eight groups of travelers based on some number of dimensions, have value today, in a world where you can do what you were just talking about? Are these just new capabilities that help us to pursue our fundamental thought process more effectively? Or are we leaving the fundamental thought process behind?

O'TOOLE: We're taking the conventional segmentation, advancing it to the individual level, and optimizing it. People have talked for years about doing this, but now it is the established practice that we continue to refine.

I'll give you a specific example. United Airlines serves and must serve and will serve a wide range of customer segments. We value and want to provide a good travel experience to all of our customers. That said, there is an identifiable customer segment—"high-value customers"—that produces a greatly disproportionate share of the revenue and the value creation. We focus intensively on driving our share of business, not just at the segment level, but at the individual level, for these customers. We monitor and score their business with us and their experience flying with us for each individual customer on each individual flight, and over time we get to know how we're performing with that customer and how we can recapture, retain, and grow their business.

So, yes, we're still talking about customer segments, but now we're focusing on the individual level within customer segments.

Another key shift is that we are now in a world of ongoing, dynamic testing and optimization. If you want to sell Economy Plus seats, do you market the benefit of "more legroom," or "more space to work," or "more space to stretch out and relax"? Before, there were conventional ways to test this, but not in real time. Now, we can and do test and optimize dynamically

Create a Data-Friendly Culture

Tip 3. Make Curiosity a Criterion for Advancement

Organizations must define intellectual curiosity as a basic criterion for advancement.

O'Toole likens the notion to the requirement that employees communicate effectively and work well with others in the organization. While a difficult skill to measure, it is unquestionably a practical and important factor in promotion decisions, particularly as one advances to higher levels.

Companies also need to provide their employees enough leeway to explore the questions that interest them—even if these questions do not present obvious applications or quick revenue prospects. In his view, the risk of overfeeding curiosity is better than the cost of stifling it.

in real time, based on specific customer attributes.

I'm quite mindful that people have talked about "a segment of one" and personalized marketing for years. That's not new. What I think is increasingly advanced is the capability to do highly individualized marketing driven by robust data analytics at very large scale and to optimize it dynamically in real time.

LEININGER: Making this shift requires different personnel and priorities. Let's talk about the people involved. You've been in your job for six years. How does your marketing team look different now than it did six years ago, and what do you think it's going to look like six years from now?

O'TOOLE: I hear people say, "We need data scientists." Well, yes, very selectively—but what you need more broadly are people in different types of functions who are able to translate business needs and problems into data analytics, manage the data required, perform the analytics, and then apply the analytic output in the execution of marketing initiatives and activities.

Create a Data-Friendly Culture

Tip 4. Demonstrate Intellectual Honesty

Curiosity paired with data can generate unexpected insights. But these insights will be worthless if they are disregarded or shut down.

"Don't reject answers just because they are inconvenient or don't support your parochial view or functional role or opinion, or because they call into question established practices—in simple terms, because they aren't what you want to hear," says O'Toole.

Even subtle criticism from senior leaders can sharply curtail people's willingness to bring forward honest information. "If you want to torpedo an intellectually curious culture very quickly, that's a good way to do it."

> The conventional marketing organization is an artifact of a prior era.
>
> **TOM O'TOOLE**

The basic, minimum level of understanding of how to use data analytics has risen, and continues to rise. For virtually any marketing manager, it's foundational to your daily job. For the CMO, proficiency in the use of data analytics is imperative.

I see companies continuing to retrofit and adapt and incrementally change conventional marketing organizations—but ultimately what is required is a restructuring of the marketing organization, which is a wrenching transformation. The conventional marketing organization is an artifact of a prior era.

LEININGER: What does that restructuring look like?

O'TOOLE: If you really strip it down, the core elements include data, content, and platforms. Content includes what was previously called "advertising" or, specifically, "creative." Platforms include what was previously called "media." Foundational to it all are data and data analytics. That's rudimentary, and certainly not the whole framework, but I think it's a rough start.

The conventional marketing functional

structure is increasingly anachronistic. The functional distinctions are now at best arbitrary and becoming problematic.

Look at e-commerce. Where does the marketing function leave off and the IT function pick up? That's a stress line in every organization I've ever been involved with. Look at data-driven advertising. Where does the data-analytic function leave off and the advertising function pick up? That is a stress line right now and will become more so.

LEININGER: Organizations created separate digital-marketing and e-commerce teams because their legacy marketing teams couldn't do what needed to be done, and if digital marketing or e-commerce was made subordinate to classically trained marketers, change wouldn't happen fast enough.

Now organizations are saying, "I created these separate teams so I could get something done—and now I have three organizations where really I need one unified organization." What do you see as the biggest challenge of getting from those siloed organizations to one organization?

Create a Data-Friendly Culture

Tip 5. Turn Information into Action Promptly

Ultimately, a company needs to take action based on data-driven insights. Though there is no single road map for this, executing promptly is crucial to success.

"Too often, companies initiate a proposal for a multiyear plan subject to funding approval and IT prioritization that will start next year to enable or act on data analytics," O'Toole says. "Sometimes that's necessary, but you need to be asking, 'What can be done right now? How can we use data insights tomorrow or today?'"

O'TOOLE: Restructuring the marketing function requires getting comfortable with and proficient at much greater integration between what were conventionally different functions. In the era of broadcast network television commercials, how much integration was required between broadcast network television advertising and consumer data analytics at the individual level? Not much. (I'm not referring to media planning or audience research). But how much integration is required between targeted digital marketing through addressable media and analytics? A lot.

LEININGER: How unique is this trend to the airline industry?

O'TOOLE: I don't think that it's at all unique to the airline industry. Yes, travel and financial services and retail are relatively advanced. But it also applies increasingly to healthcare and energy and a wide range of other industries. I think that it ultimately will apply and can apply to virtually all industries.

I was giving a presentation once about exactly this subject: data-driven marketing. A guy came up to me afterwards and basically said, "Yeah, that all sounds great, but I work for a company that makes steel. We're a commodity, so how does this apply to my company?" I don't know much about the steel business, but I thought about it and said, "Well, do you make different kinds of steel? I imagine you make rolled steel, formed steel, cut steel, tooled steel." He said, "Yeah." I said, "I imagine you have different types of customers and that you know who buys what type of steel. Couldn't you get the data and figure out which customers might be prospects for specific types of additional products or offers?" And he said, "I suppose that we could." I think that it's really a matter of reframing the business and looking at it in terms of what data is available and how it could be used.

LEININGER: How do you use the richness of the data that you collect to reinforce loyalty or a sense of relationship with a customer?

O'TOOLE: It lets us track and evaluate the customer experience over time and intervene if we see the relationship weakening or how to grow it. It enables us to be very, very informed about individual customer relationships.

I'll give you an example. The most valued benefit of being a Premier member of MileagePlus for many people is the complimentary upgrade. The availability of complimentary upgrades on a given flight is a function of the bookings on that particular flight and the resultant seat availability in the front cabin. It sometimes happens, just by circumstances, that a 1K member can have an extended series of flights without a complimentary upgrade. This can put us at increasing risk of losing a high-value customer.

Now we track if an individual has had a run of X number of flights in a row and their upgrade never cleared. We can then automatically override the algorithm and intervene to be sure that the individual gets upgraded. This has been proven to result in a measurable business gain.

LEININGER: You've got all these smart people on your marketing team, and they're trying to figure out how to get more cash from each transaction. How do you make sure that, as they write those algorithms to maximize cash, they do so in a way that is consistent with how you want the brand to be perceived?

O'TOOLE: In terms of the brand, there are certain things that we could do that we just decide not to do based on brand considerations.

I'll tell you a different and growing question that we think about: operational complexity. Here's an example. One of the things that we know is potentially popular with travelers is offering preordered meals, where instead of getting three choices, here's a menu and you can choose from 25 choices in advance. That all sounds fine until the flight attendant needs to know that the guy in 12B had the kale and quinoa. And what happens when there is bad weather, and to get you out sooner, we swap aircraft or put you on a different flight, when you paid in advance for the kale and quinoa?

The point is that we can, through digital channels and data use, create highly differentiated service offerings that can cause enormous operational complexity at scale in the real world. Thus, we increasingly need to balance service individualization with operational execution to find the optimal combination of personalized products and services with reliable operational service delivery in a highly dynamic business on thousands of flights around the world every day.

Featured Faculty

Tom O'Toole, Clinical Professor of Marketing at Kellogg

Eric Leininger, Clinical Professor of Executive Education at Kellogg

Related Resources

Read: 5 Ways to Attract and Retain Data Scientists

Read: How Marketing Leaders Can Thrive Amidst "Tectonic" Shifts in Expectations

Learn more from Tom O'Toole and Eric Leininger in our **Executive Education programs**.

How to Build (and Keep) an Analytics Team in an Established Organization

A look at how Vanguard attracts and empowers a community of data scientists.

FOR DECADES, THE PROPOSITION was simple: if individuals wanted to invest, and to do so affordably, they turned to Vanguard, a pioneer in passive funds tied to large stock-market indices such as the S&P 500.

Over time, the popularity of index-fund investing has skyrocketed, and today Vanguard is the largest provider of mutual funds in the world.

But all this success has come with some unexpected consequences for the firm's strategy for reaching clients, says Jing Wang, Head of the Center of Analytics and Insights at Vanguard. "The low-cost index fund—it used to be the underdog. It took us thirty years, but now it's mainstream. Now everybody is copying," she says.

This, combined with other trends, such as ever-increasing expectations around the customer experience, prompted Vanguard to rethink how it finds and serves its clients. So three years ago, Vanguard gave Wang a new mandate: develop the company's consumer-analytics capabilities from the ground up.

Wang recently sat down with Insight to discuss the team she built: how it is structured, her strategy for hiring and retaining data scientists, and how her team gained a reputation for solving some of the company's thorniest problems. She also explains why having support from senior leadership is critical—but no substitute for the power of actually *showing* people what analytics and AI can do.

This interview has been edited for length and clarity.

KELLOGG INSIGHT: Please start off by explaining what, exactly, your team has been charged with doing.

JING WANG: My team is called the Center of Analytics and Insights. We have three functions. One is a market-research function that I inherited from the previous organization. The next is our data-science and machine-learning group. And then the third is the analytic-enablement function, which we built. It's a small function, but it primarily focuses on a few things, including analytics training and education at every level of the organization, and pulling together analytics practitioners from different functions to build a community.

INSIGHT: Tell us more about that final point about building a community around analytics. I would imagine at a company like Vanguard, there is data science everywhere. There must be data scientists working throughout the organization on a huge range of problems, some of which probably have nothing to do with customer analytics. So can you talk about the broader data-science ecosystem a little bit?

WANG: When you say "data science," it's not a term that people have uniform understanding of. Anybody who can manipulate data, try to generate insights from data, produce models, build AI capabilities—they could all be data scientists, right?

So to answer your question, Vanguard has traditionally had very strong analytic capabilities in the investment area. But they don't call themselves data scientists. They call themselves investment analysts or financial analysts. So part of the analytics-enablement work is to get everybody together and agreeing on a consistent set of levels and job grades and to align those with skill sets and competencies

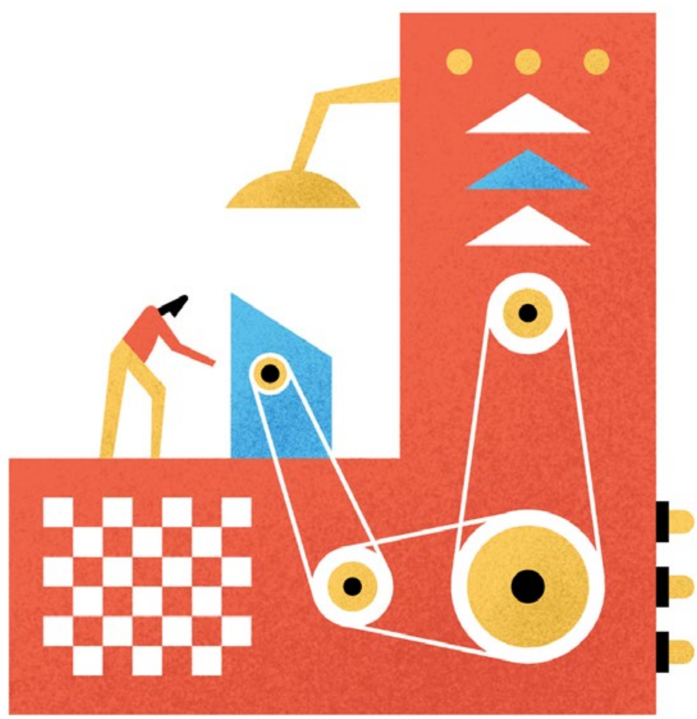

so someone can actually say, "Hey, I'm working in this division. I'm doing all these things at a level two. So here's what I need to do to get to level three in another division, another group." We also have SharePoint sides for people to share best practices or post jobs. This helps to cultivate a sense of belonging for the analysts. It also helps them help each other to resolve challenges.

INSIGHT: So you would definitely consider it a win if there were data scientists working in your division, doing client work, who took their next role over with the investment team? You would love to see more fluidity in terms of a career path in data science throughout the organization?

WANG: Well, yeah. I think that's pretty critical to retain those kinds of talents, right? Because if you don't help a person to understand how his career path could keep him engaged and growing, I don't think you're going to be able to keep this person. It's also a win–win. We realize that we can help to transform the company culture by playing an active role in bringing in new talent, incubating them here, and looking for the right opportunities to send them to other business divisions.

INSIGHT: What was the impetus behind creating this Center for Analytics and Insights three years ago? What about the changing landscape made Vanguard say, hey, we need to rethink how we are serving our customers and clients?

WANG: Vanguard has been incredibly successful in terms of serving investors. Our whole purpose is to really put a client in the center—to take a stand for all investors, to treat them fairly and give them the best chance for success.

In the past, our strategy had been based on providing a low-cost, broadly diversified investment product in a very high-cost overall investing environment. I think this—combined with Vanguard's fundamental corporate structure as the only investment-management company owned by our clients, meaning there's no conflict of interest—allowed us to build a very strong reputation as trustworthy. For example, after the global financial crisis, many investors actually turned to Vanguard, because we are trusted.

So that strategy has made us successful, but we also came to realize that the same thing that made us successful may not continue to make us successful in the future. The low-cost index fund—it used to be the underdog. It took us thirty years, but now it's mainstream. Now everybody is copying. So that position becomes hard to articulate.

And the client-experience factor plays such a huge role in a digital environment. The clients have been trained by the Amazons, the Apples of the world and have much higher expectations about their interactions with Vanguard. In the past they've been willing to say, "Hey, we know Vanguard doesn't have the greatest website, but they're a company we trust, they have a great product." But in the future, those kinds of experiences are going to really shift and impact client loyalty.

This led us to believe that we need to become better at using analytics to provide a much more individualized experience for our clients. We serve millions of investors directly. We know exactly what they do. And we have not done enough to use that insight to engage with them in a more personalized way to deliver greater financial and emotional value from investing with Vanguard. We need to treat that information about their needs and preferences as an asset.

INSIGHT: What have you learned in the process of designing your center?

WANG: Implementation is key. You can build a beautiful model, but if it is not implemented, it's useless.

When I think of implementation, I start to think about automation. So, for instance, instead of feeding you raw data for you to digest, what if we just suggest three potential clients every day and say, "Go talk to them." But that takes a lot of change, right? That takes technology, that takes behavioral change, that takes a mindset shift. It takes, to be honest, a lot of leaders starting to think differently and manage differently as well.

INSIGHT: You had a pretty strong mandate from the CMO to make these changes. So presumably you felt a lot of support from the top of the organization. But creating this ecosystem involved pretty widespread change *throughout* the organization. What kind of resistance did you face in trying to change the way clients were served?

WANG: Our first challenge was getting businesses to come to us with questions. At the beginning, people would come to us thinking we're a reporting function, that we can manipulate data. And they would tell us, "I need A, B, C, D." We would have to step back and say, "What are the business questions that we're trying to solve? What are our objectives and goals?" And we'd come back with a solution that brings different elements and components together. It may not be the components they had in mind. It may be the case that a much bigger or smaller solution actually addresses their initial question. So initially, it takes a lot of education to show the possibilities. Now people know we're here. We have this reputation for solving the most complex questions for the business.

The second challenge is that business leaders tend to have developed certain beliefs, based on their experiences, that might not be correct in every instance. So we still have to show the business that their hypothesis is not the right hypothesis. We find some resistance if the answer we bring them is different from what they expect.

It's interesting that the B2C and B2B areas are different. I don't face a lot of resistance at a front-line level on the B2C side, because we have so many individual clients. If you tell people working in B2C that using this tool can allow them to have a really great conversation with the client when they call in, that's great. It makes their job easier.

On the B2B side, execution is delivered through a human: the relationship manager, the salespeople, some of whom have been very successful for twenty years in their territories. And it can be harder to encourage them to change their process.

INSIGHT: Implementing these kinds of difficult changes must require the right team. Do you have anything in particular that you look for in terms of hiring and team composition?

WANG: Yeah, there are a couple things. When we started, we had a small, uniform analytics team. Most of them were skilled practitioners in the marketing-analytics area. And we realized that we needed to build a team that has a much wider range of training and expertise. We needed to have senior-level data scientists, we needed to find a junior-level data scientist, we needed data engineers, we needed engagement leads—you know, "translators," because it's a lot of effort to translate analytics into practical business applications. They all needed to work together. And we also added, this year, machine-learning engineers.

Even within data scientists, we need different skill sets. We need people who are really good at natural language processing. We need people who are really good at deep-learning optimization. We need people who are really good at statistics, et cetera. So people start to have sort of different flavors and expertise. And we also love people from different backgrounds. We have people with a PhD in physics, a PhD in math, a PhD in engineering. We are going make an effort to hire a PhD in social science. Those different perspectives help to open people's eyes to different ideas and create innovative solutions much faster than before.

INSIGHT: Where do you find all the people you need?

WANG: There's no one way; this is something we have to always work hard on.

We love having a dedicated HR person. Having someone who actually knows exactly what we do and, roughly speaking, what the relevant experiences and skill sets are has made the hiring process much easier.

We also build a number of relationships with training programs, like the Insight Data Science Fellows program, that become a really good feeder of entry-level talent. And we have started to build relationships with local universities. We attract the interest of students, so we can bring them in as interns.

INSIGHT: What advice would you have for the version of you three years ago? What knowledge would you want to give somebody trying to start up an analytics capability in a large organization?

WANG: It really helped me that from very early on I had a very strong sponsor. He is one of our senior staff, and he reports directly to the CEO.

Him being personally interested in investing in the business area has been helpful, because it has a lot of signaling value.

I am not a data scientist, so I also found it was effective to hire a terrific data-science leader with a strong growth mindset who was able to partner with me.

I think the final thing I will say is that, although I knew it was going to be difficult to address that data-ecosystem problem, I *still* could not have prepared myself for the amount of change management that we did.

Initially, I thought we just needed leadership support. But I came to recognize that that alone was not going to solve the problem, because senior leadership cannot magically snap their fingers and everything will happen. We really have to work with all the layers of the organization, including the IT team that actually built the application.

And that is hard, because you have limited resources and energy. So where do you focus your effort? We constantly asked, "What are a few big things where, if we make a difference there, it will help open people's eyes to see what consumer analytics is about?"

You have to *show* people how it's done. Because it has never been done before.

Featured Expert

Jing Wang, Head of Center for Analytics and Insights at Vanguard

Tips & Tools

Based on insights from **Steve Franconeri**

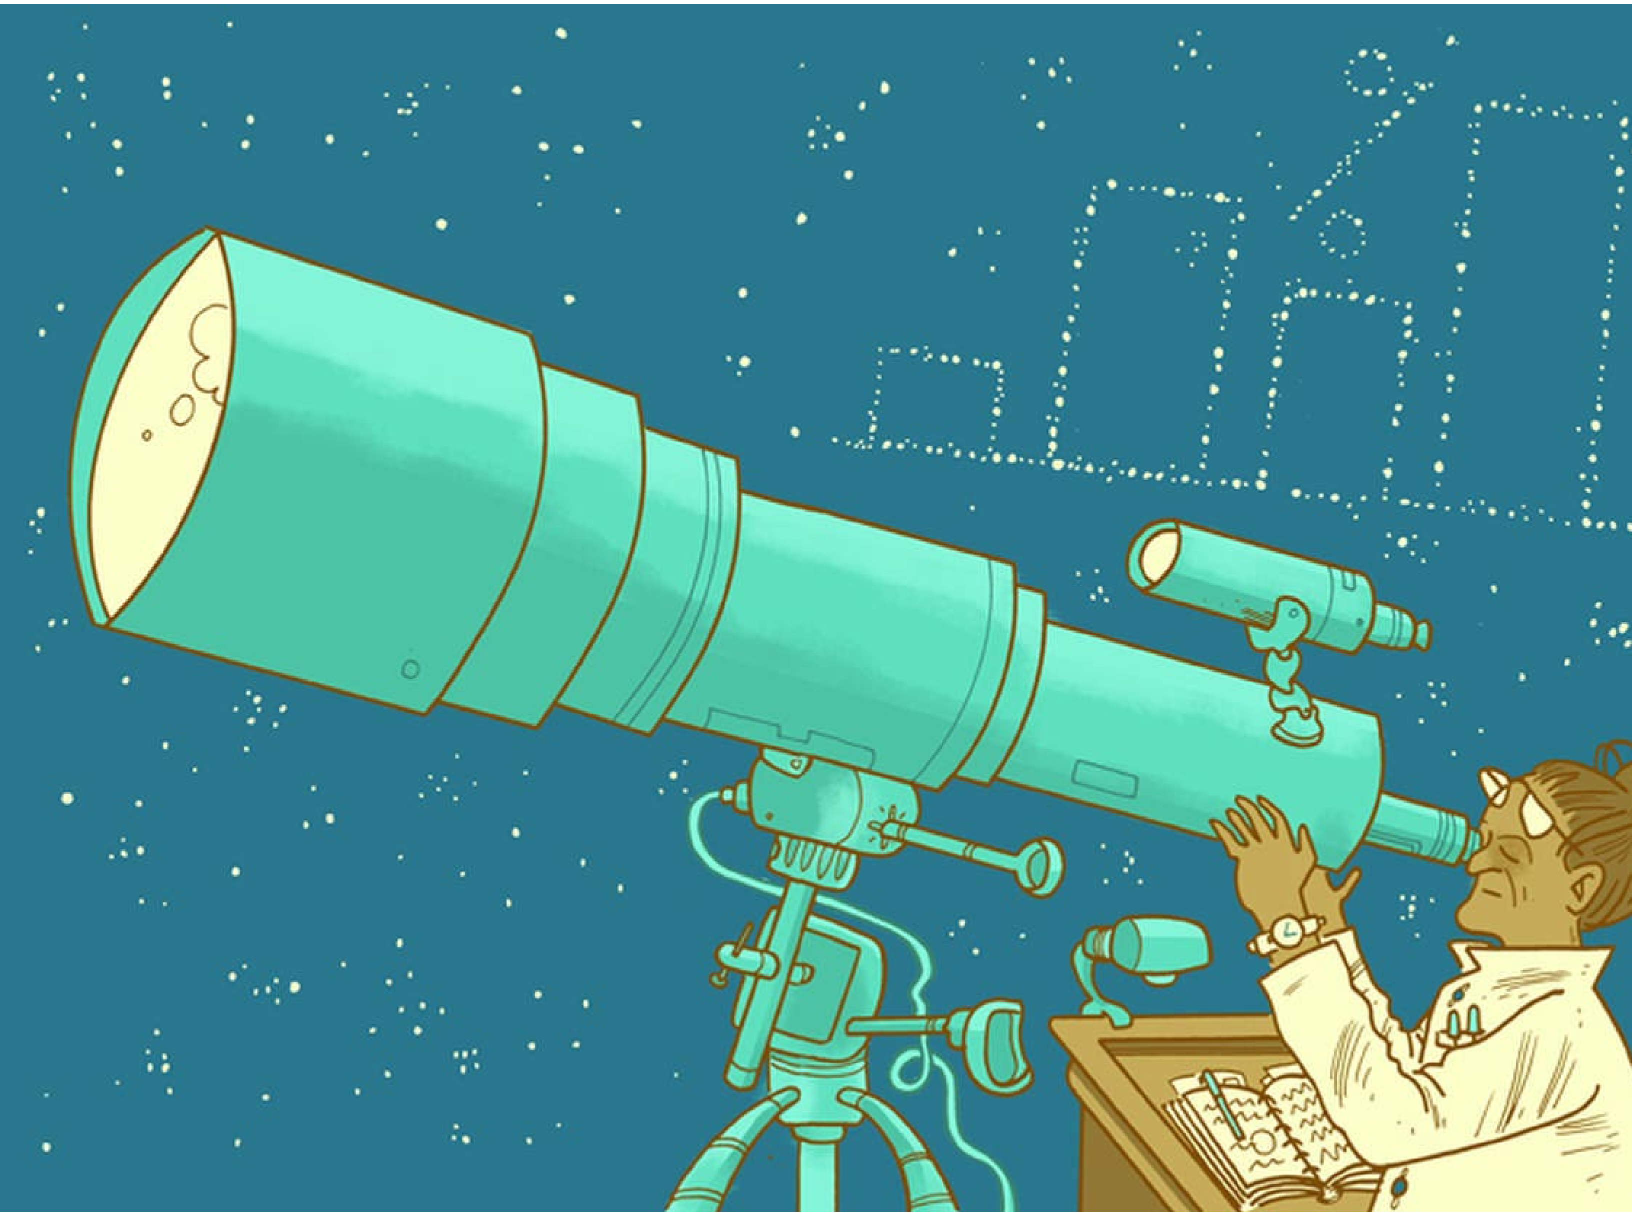

How to Tell Persuasive Stories through Better Data Visualizations

Nine rules for mastering this essential skill.

YOU HAVE A GREAT IDEA TO GROW your business, and it's based on solid data. How do you convince your skeptical boss, staff, or clients of your brilliance?

Show them your data visually, says Northwestern professor Steven Franconeri. "People might think that visualizations are pretty and they're icing," he says. "That's not true. They are indispensable, absolutely indispensable."

Around 40 percent of our brains are devoted to visual comprehension, according to Franconeri. "It's critical that you use that machine."

Data visualization has been around for more than 100 years. Florence Nightingale was an early pioneer who used a visualization to prove that more soldiers died while lingering in unsanitary hospitals than on the battlefield.

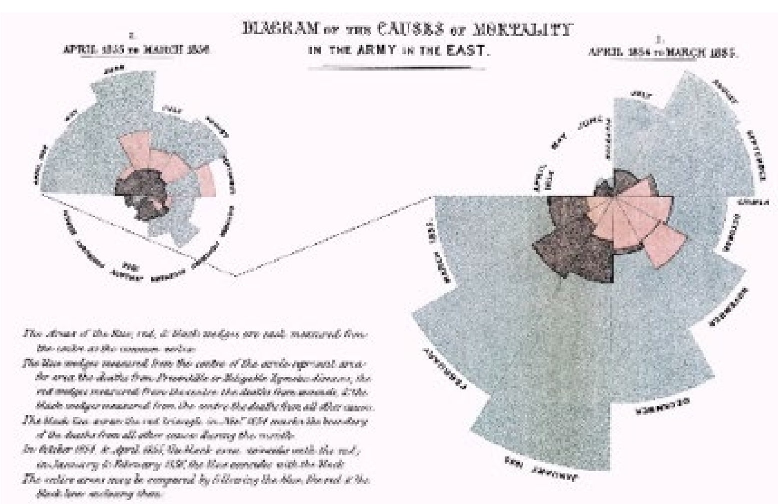

This classic data-visualization example was created by Florence Nightingale to show that more soldiers were dying from preventable diseases after battle (blue) than had died from wounds on the battlefield (red).

Credit: Wikicommons

But in the past decade or two, data visualization has become more sophisticated. The technology for displaying data has improved, and research is establishing best practices for communicating data visually.

This boom has led to some rules about how to best present your data. Yet Franconeri, who also has a courtesy appointment at Kellogg, says far too few people learn those rules. And it comes at a cost: bad visualizations.

"If you do it wrong, it's a disaster," he says. "It can be complete spaghetti."

Boom in Visualizations

The way you display your data depends on the story you want that data to tell. So it follows that the people at the forefront of data visualization today are trained storytellers.

"The folks who are doing the most interesting work on the storytelling front are journalists," Franconeri says. They are the people whose job it is to take complex ideas and distill them in digestible, memorable ways.

Take, for example, this interactive graphic that looks at how the recession shaped the economy and this one that looks at how the yield curve predicts economic growth. In both, from *The New York Times*, the viewer is guided through several sets of related data with an animation that causes one visualization to flow into another (either by scrolling or clicking).

"It's taking a big complex data set and guiding people through one snapshot at a time so that they understand the big picture through a guided tour," Franconeri says.

And all the rules that apply to good journalistic visualizations also apply to the business world.

"Nobody ever gets taught these rules. You take writing classes in college. You don't take a graphical-communication class," he says. Yet, "this is

> If you learn these rules, it will have a multiplicative impact on how well you can convey your ideas to people and how well those ideas will actually sink in and then lead to action.
>
> **STEVE FRANCONERI**

a skill that people need to have. If you learn these rules, it will have a multiplicative impact on how well you can convey your ideas to people and how well those ideas will actually sink in and then lead to action."

Here Franconeri offers up a few of the most important rules.

1. Replace text with visuals in oral presentations.

When Franconeri teaches business students and leaders, he advises them to step away from the traditional bullet-pointed slides and start integrating visuals into their presentations.

Because our brains use the same systems to process speech and written language, putting up text on a screen while talking "ensures that you won't get your point across, because no one can read and listen at the same time," he says.

But, he says, "you can look at pictures and listen at the same time at full blast."

2. Visualizations can make it easier to make sense of information.

Visualizations allow our brains to process dozens to thousands of numbers simultaneously.

72	50	64	18	23	47	48	67
64	87	61	18	40	64	88	15
17	31	78	41	56	87	42	33
63	37	59	15	11	55	30	63
55	48	49	22	22	56	66	61
48	21	79	34	45	84	57	50
21	94	54	12	28	20	94	19
80	37	50	64	54	31	38	74

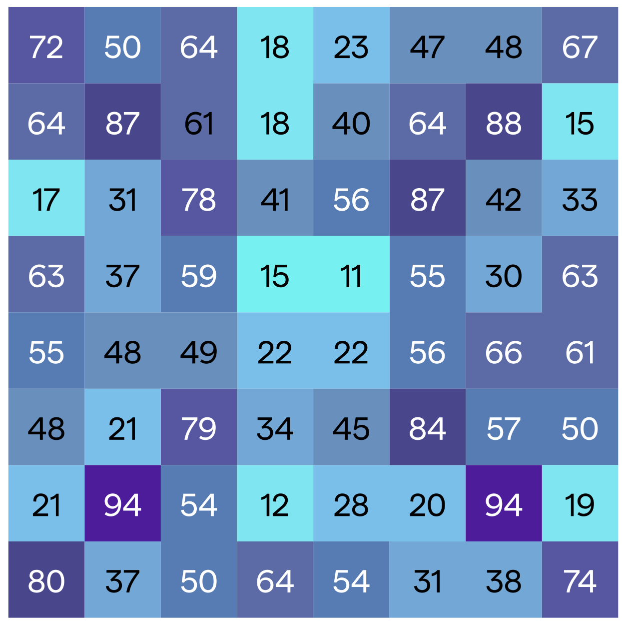

Credit: Steven Franconeri

Look at the first grid (with no formatting), and try to pick out whether there are any consistent patterns in the numbers. Now look at the second grid, where the data are visualized through different colors, and see the pattern of high values in the diagonals that you previously missed.

3. Visualizations are uniquely persuasive.

Research on visualizations supports the idea that they can be perversely persuasive. This is because we are much more likely to believe conclusions if they're visualized than if they are just given in numbers.

Franconeri describes a study where participants were told about a trend in data—either a military-troop surge, a change in unemployment numbers, or a trend in climate-change data—with verbal descriptions of the change in these numbers. "Then you give the people the information in one of two forms: just that text or text accompanied by a graph," he says. People were much more likely to report believing the trends in the graphical formats.

"Something happens in your brain when you get the visualization; information sinks in more deeply and is far more difficult to ignore," says Franconeri.

4. Visualizations are powerful—but they can also be powerfully misleading.

Still, visualizations do require time and effort to understand.

"When you look at a picture, you recognize objects, faces, or Pokémon within 150 milliseconds. That's *half* the duration of your *fastest* eyeblink. Visualizations don't work like that—they need to be *read*. Understanding patterns in data isn't like recognizing a picture; it's like reading a paragraph," says Franconeri.

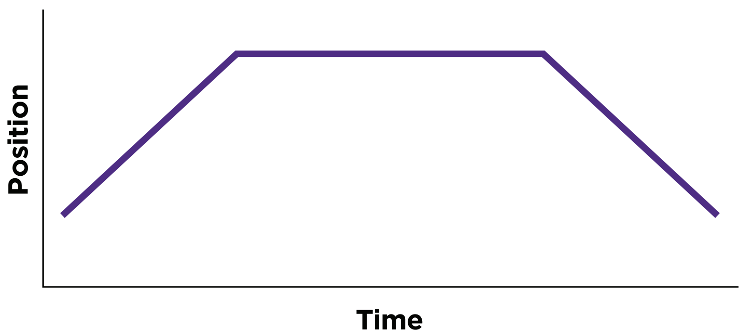

This image, based on one used in Kozhevnikov, Hegarty, & Mayer (2002), offers a good example of how a seemingly simple graphic can be "misread" and lead audiences astray. What does this graph show? Can you describe it in a single sentence?

Franconeri explains that many people, even highly educated college students, say that it clearly depicts a person walking up and down a hill, or escalators in a mall, or perhaps an airplane taking off and landing.

"That's treating the graph like a picture from a camera. It's using the part of your brain that is practiced in seeing and understanding the world outside," he says.

What this graph is actually depicting is an object that starts off moving (in any direction), then stops moving (because time continues changing, but the position is not changing), then moves again.

"A visualization is not a 50-millisecond view, where you look at it and you understand what's happening. It's a minute-long process," says Franconeri.

5. The way that you organize the data controls the pattern that the eye will see.

If visualizations need to be read, what controls how you will order each of the "sentences"?

Take the pair of graphs below, which display shares of aggregate income by type among individuals aged 65 and over in the U.S.

The graphs show the same data. But because the data are organized differently, people who encounter the graphs will likely approach them differently and thus may draw different conclusions.

For the graph on the left, people tend to compare income types within each year—noting, for instance, that Social Security income was higher than earnings or government pensions in 1962. For the graph on the right, however, people tend to compare values across years for each income type—noting that most income types have increased over time.

“The visual system can’t compare every value to every other value. That’s too complex even for such a powerful system,” Franconeri explains. “So it tends to prioritize comparisons of values that are close to one another, and you can use that rule to control the comparisons that you make when you are analyzing data or that your audience makes when you present.”

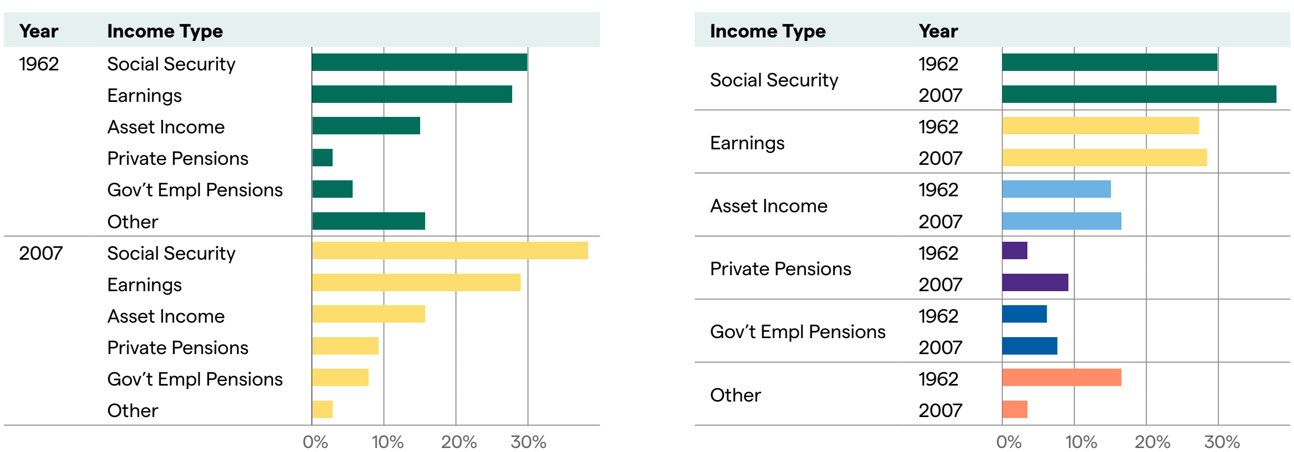

Credit: Steve Franconeri. Data from the Social Security Administration.

Interestingly, the same principle applies for word clouds—an increasingly common way of displaying textual data. It's tough to find the common themes on the left, but it's an easy task when you clump up the related words for easy comparison.

Credit: Marti Hearst

6. Give your visualization a test run. Or three.

Still, there are a lot of factors that go into how we process visualizations. Even professional designers can't always predict how a visualization will be read. This is why it's critical to use test runs.

If you have a relatively simple data set, create three or four different visualizations. "Make them as different as possible," he says. "Show those to some colleagues. Ask them what story they see in each one. You'll be surprised at the differences in what people extract from the same data, plotted in different ways." Those interpretations can then be used to shape your final product.

With a more complex data set, keep in mind that your viewers can quickly get overwhelmed even if you and your designer understand the visualization completely. Because the visual-processing centers in our brains are so strong, it is easy for designers to feel completely comfortable with a complex visualization because they have been immersed in the data. Viewers, however, struggle without that background.

"Their brain is lit, but they don't know what to do with it all," Franconeri says.

7. To convey more complexity, consistency is key.

By the same token, you have more leeway to create complex visualizations when you show people the same type of graph again and again.

"People can handle an impressive amount of complexity once they get used to it. People gain expertise in moving their eyes and attention around in a certain pattern, looking at the shape of this and the slope of that. There's a strong visual expertise that gets developed," he says.

"Think of the flight deck of an airplane. It's an insanely complex set of displays and controls. With practice, pilots become completely comfortable and the complexity fades away. But that expertise is incredibly specific: change the layout of those displays and controls, and your flight is in serious trouble."

8. Innovative formats *can* be attention-grabbing and effective. Just make sure you provide your audience with support.

Visual storytelling needn't be stodgy or predictable. Innovation is possible, so long as you tread thoughtfully.

Take, for instance, the connected scatterplot—a technique whose effectiveness Franconeri has researched. It takes a style of graph that is familiar to most of us—a simple scatterplot displaying two different trends—and merges the two trends into a single visualization, as in the purple line below.

In a traditional scatterplot graph, time would run along the x-axis, with the y-axis showing two different sets of data. In the example below, these data sets are miles driven by Americans each year, as well as auto fatalities each year.

Instead, a connected scatterplot combines the data from the two y-axes and lets time run as an annotated line across the graph. When Americans on occasion reduce their driving from one year to the next, the line goes backwards.

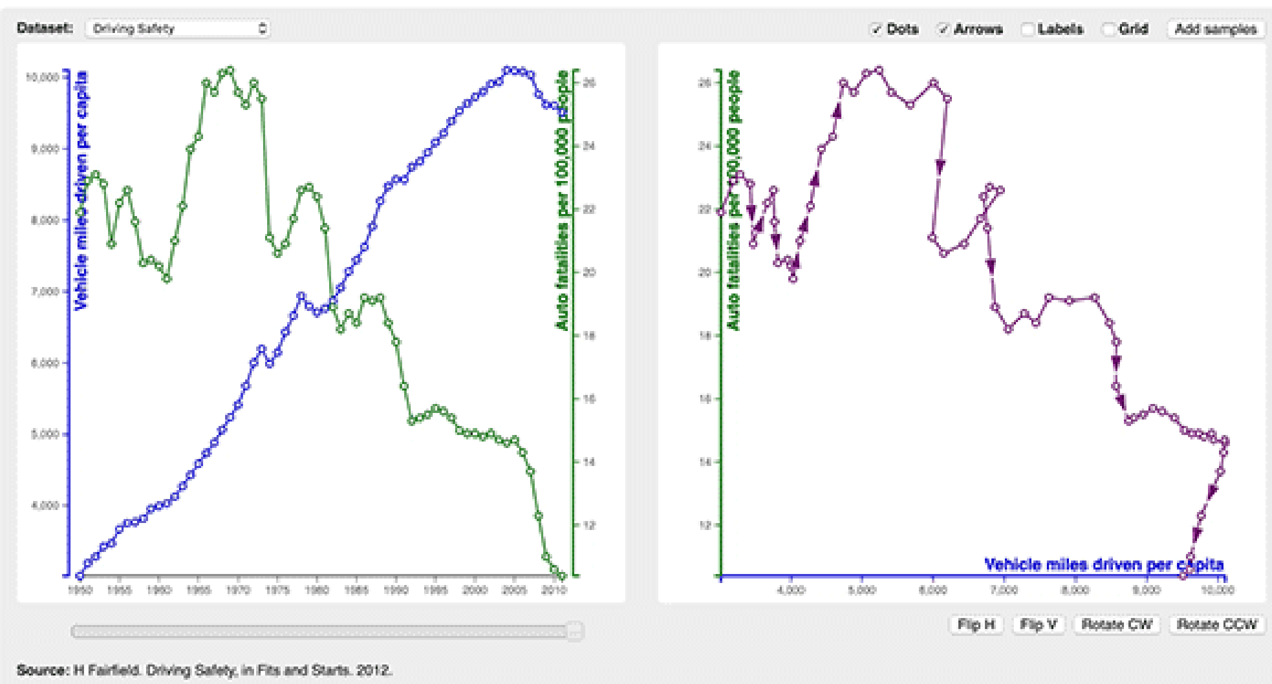

Credit: Steve Haroz. This scatterplot is derived from one originally constructed by Hannah Fairfield for *The New York Times*. To interact with the connected scatterplot yourself, visit Haroz's site here.

Most viewers do not expect a line on a graph to behave that way.

Yet Franconeri's research suggests that novel formats like this can be highly effective at grabbing viewers' attention, producing what education researchers call a "desirable difficulty" that engages viewers with a puzzle to be solved. And with a clear explanation, most people can understand the data quite well under these novel formats. But this is where those test drives become especially critical: wherever experience suggests that people get tripped up, you would want to consider adding clear annotations, as well as the option to view more traditional visualizations of the data in specific spots.

9. Sometimes visualization isn't the answer.

Visualization is a wonderful tool, but it is not the only one at your disposal.

Our visual system is not necessarily equipped to make sense of data with more than two or three dimensions, Franconeri explains. So when you get into more complex models with many factors and drivers, you may have to rely on equations or algorithms to communicate your data. (And sometimes there is a fine line between successfully delivering on an ambitious infographic and overreaching. Recall that infographic about the yield curve published by *The New York Times*. "Even very sharp readers struggle to hang on when understanding patterns that are this intricate," says Franconeri. "This was a really bold attempt.")

Another time to look beyond visualization is when the data are too abstract. Instead of trying to visualize the idea that "customers in Demographic Group X of Age Y and Income Z are not interested in category A," Franconeri explains, it is often far more effective to engage the parts of your audience's brain that think about real people in real situations by considering a concrete person from that demographic: "Here's Mark, and he rides a motorcycle, and he's not interested in category A. What do we do?"

Featured Faculty

Steve Franconeri, Professor of Psychology at Northwestern University and Professor of Management and Organizations at Kellogg (Courtesy)

Learn more from Steve Franconeri in our **Executive Education programs**.

For Smarter Analytics, You Need to "Purposely Mess Stuff Up"

Why the best firms are purposely injecting variability into their data.

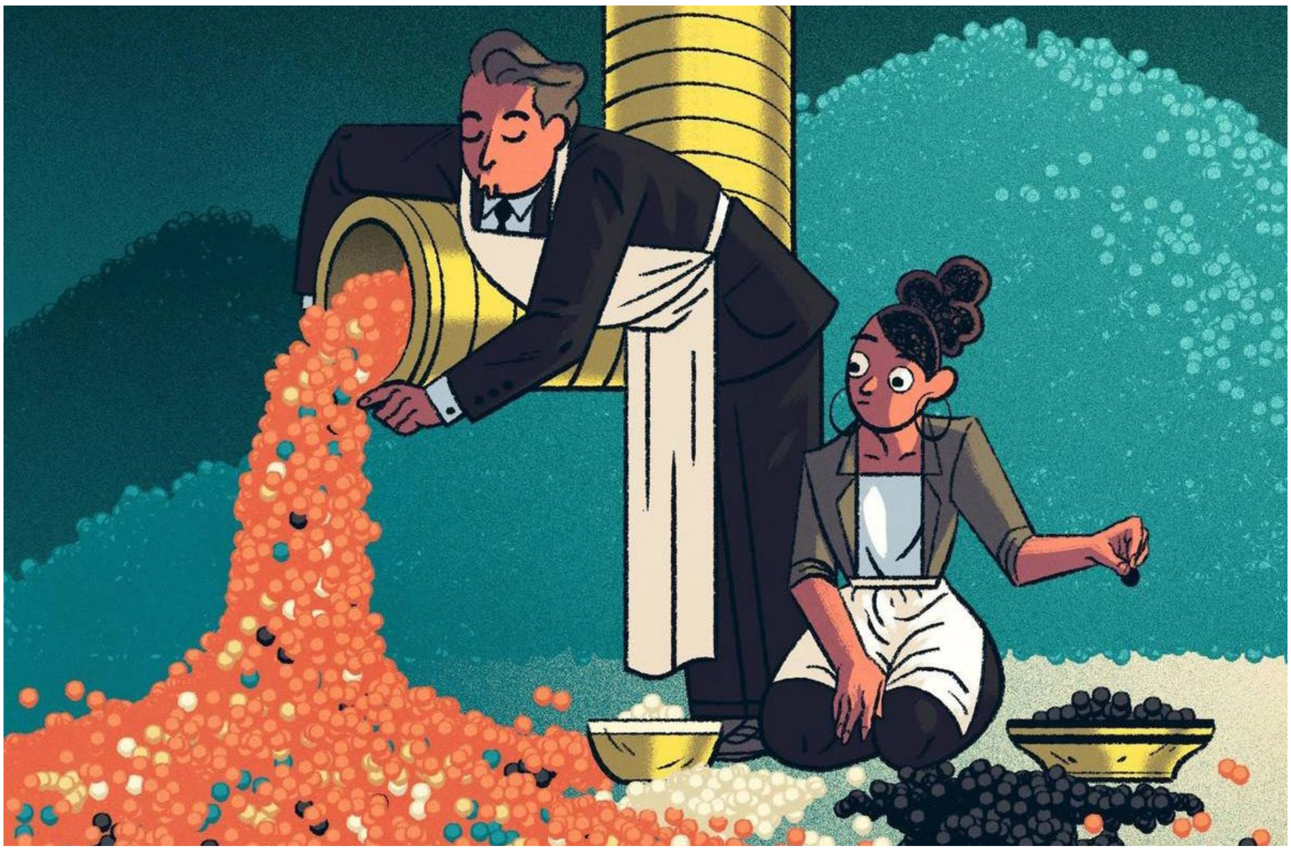

THERE'S A PARABLE THAT ERIC ANDERSON, a professor of marketing at the Kellogg School, likes to tell, one he's deemed the "Analytics Paradox."

"A young firm starts out making many mistakes. Eager to improve, they collect lots of data and build cool new models," he says. "Over time, these models allow the young firm to find the best answers and implement these with great precision. The young firm becomes a mature firm that is great at analytics. Then one day the models stop working. Mistakes that fueled the models are now gone and the analytic models are starved."

The paradox is that the better the firm gets at gleaning insights from analytics—and acting on those insights—the more streamlined their operations become. This in turn makes the data resulting from those operations more homogeneous. But over time, homogeneity becomes a problem: variable data—and, yes, mistakes—allow algorithms to continue to learn and optimize. As the variability in the new data shrinks, the algorithms don't have much to work with anymore.

The paradox leads to a rather startling recommendation: "Occasionally, you need to purposely mess stuff up," Florian Zettelmeyer says.

"You design variation into your data in order to be able to derive long-run insight," explains Zettelmeyer, also a professor of marketing at Kellogg.

Zettelmeyer and Anderson are academic directors of Kellogg's Executive Education program on Leading with Big Data and Analytics; they are also writing a book about data science for leaders.

Here, they offer a look at how the best firms have found a way to sidestep the Analytics Paradox.

From Optimization to Stagnation

In some sense, the value in big data lies in its messiness—in the often unexpected variation in how events play out and the myriad ways these events help establish connections between variables that can help people make better decisions.

"In theory, the best manager for analytics is the one who walks into the office every morning and flips a coin to make all decisions," Anderson says. "Because if you make all your decisions by flipping a coin, you will generate the best possible data for your analytics engine."

"The problem," he adds, "is that at every company, the manager flipping coins gets fired very quickly. The managers who survive are the ones who are really good at implementing decisions with great precision."

To understand how the best teams can find their operations too optimized for their own good, Anderson offers this hypothetical example.

"Right now your company offers two-day delivery, and someone says to you, 'I would like you to go back and analyze the historical data. Tell me whether we should have two-day delivery or move to one-day delivery.' Could you answer that question with your data?"

If your delivery process is being overseen by a high-performing team focused squarely on efficiency, then you likely *cannot* answer this question with data.

"If you are really good at delivery—if you've been running operations efficiently—how many days does it take? Two days," says Anderson. "The guy who was messing up and taking four days to deliver a package was fired. The one who was delivering in three days sometimes and one day other times got fired. You're left with all of the managers who deliver in two days—you've built an organization that is so good at delivering things that it almost always happens in two days."

Hamstrung by your own success, you do not have the data to know whether a better possible delivery strategy exists, or how you might successfully move to a new model.

"If I don't occasionally do the wrong thing, I will never know whether what I think is the best actually still is the best," says Zettelmeyer.

> If I don't occasionally do the wrong thing, I will never know whether what I think is the best actually still is the best.
>
> **FLORIAN ZETTELMEYER**

What the Best Firms Are Doing

Of course, firms have plenty of good reasons for not wanting to richly reward incompetence, or promote a manager whose decision-making seems limited to coin flips.

Instead, top firms have adopted a fundamentally different strategy for thinking about big data.

"The best firms are heavily investing now in creating data, designing data," says Anderson. "They're purposefully injecting variability in the data."

Whether they are experimenting with how many days it takes to deliver a package, how to set prices, or how to best maintain an aging fleet of vehicles, these elite firms understand that experimentation and variability need to be built into the organization's DNA.

"It's just a miniscule fraction of firms" doing this, maybe five percent, says Anderson.

So what do most managers need to do differently?

"When you take a business action, you need to keep in mind what the effect is on the usefulness of the data that are going to emerge from it," says Zettelmeyer.

That requires the foresight to understand the questions you may wish to answer in the future, as well as the discipline to work backward from those questions to ensure that you set yourself up to get data that are rich and helpful.

A company rolling out a national advertising campaign, for instance, might decide to tweak the campaign in important ways only in select markets or to stagger the rollout by region. While there may be short-term costs in terms of efficiency and optimization, the resulting data have the potential to teach the company going forward.

Don't Relegate Data Science to the Data Scientists

Such foresight cannot be the purview of a single employee or team at an organization, the pair stress. That's because decisions about how to experiment should be made with specific problems in mind.

"It cuts across the whole organization, so it has to be a cultural change in how we think about our day-to-day operations," says Anderson.

The key, Zettelmeyer says, is "to transport yourself into the situation you're going to find yourself in in the future." What data would be helpful to have in order to make the *next* decision, and the *next* one? What relationship between variables do you want to demonstrate? And how could you design an experiment to demonstrate that link, given your existing capabilities and constraints?

And keep in mind that the infrastructure this requires may be quite different from what is necessary for managing much of the big data that flows through an organization. For instance, the high-level dashboards that senior leaders are used to may not be capable of distinguishing among many subtle but important differences in when a campaign was rolled out, for instance, or how a delivery route was established.

“It’s a very different thought process in terms of how you would actually build an IT system to support experimentation,” says Anderson.

Thus, rather than try to outsource this work to a dedicated data-science team—or worse, a single piece of software—Anderson and Zettelmeyer recommend that firms train managers on how to think and ask questions about data.

“It requires a working knowledge of data science,” says Zettelmeyer. “This is a skill set that managers need in order to even be conscious that this is something they need to take charge of.”

Featured Faculty

Eric T. Anderson, Professor of Marketing at Kellogg

Florian Zettelmeyer, Nancy L. Ertle Professor of Marketing at Kellogg

Learn more from Eric Anderson and Florian Zettelmeyer in our **Executive Education programs**.

Based on insights from **Joel K. Shapiro**

To Improve Customer Experience, Embrace the Outliers in Your Data

Don't ignore them. Let them highlight your biggest failures and juiciest opportunities.

COMPANIES ARE INCREASINGLY TURNING TO data science to better understand how customers interact with their products and services. And with good reason. A 2018 survey sponsored by SAS, Accenture, and Intel found that 58 percent of responding business leaders said that using customer analytics significantly increased customer retention and loyalty, and nearly half associated analytics with significant revenue growth.

"It's important to know if you're providing the experience you want customers to have and making it better," says Joel K. Shapiro, a clinical associate professor of data analytics at Kellogg. "To do that, you need to measure as much as you can about what happens to them."

But according to Shapiro, far too many companies are ignoring some of the juiciest data around. These are the data that induce head-scratching, the measurements that don't quite fit the existing models. These are the outliers. And they can highlight your product's or service's greatest weaknesses—as well as where it has the potential to truly shine.

Companies can and should use this knowledge to optimize the customer experience.

"The mere presence of outliers in customer-experience data means that really good or bad things can happen to customers," says Shapiro. "Maybe you can move that [experience] toward something that either increases the number of positive experiences or doesn't detract from them."

Keep Your Outliers

When data scientists come across an outlier, their first inclination may be to discard it in favor of "cleaning" or "smoothing out" the data. After all, the data might have been entered incorrectly or appear as the result of a modeling error. Or it may represent a freak accident—a set of circumstances unlikely to replicate itself. Why waste time accounting for the easily discountable?

Resist that urge, Shapiro says. It is always worth examining why the outlier occurred.

Shapiro comes to this insight firsthand. While studying the effectiveness of Spanish language instruction in West Virginia's middle schools, he found that overall, students who received online instruction scored about the same as those who received it in person. But he also found that face-to-face class instruction produced a wider range of experiences. Investigating further, Shapiro learned that most of the high-performing in-person students came from one specific class.

After he had presented his results, the state's board of education advised him to exclude that class from analysis because its instructor was known to work well with students. "They said she was the best thing that's ever happened to the state of West Virginia from an education perspective," Shapiro says.

But Shapiro realized that, rather than excluding this instructor's students, the board of education should do the exact opposite. They should find out what this person was doing and try to replicate it.

What this outlier ultimately demonstrated was that, under the right circumstances, the in-person classroom model had the potential to outperform online instruction. This is an important first step toward understanding the potential of these outliers, Shapiro says. "They shouldn't just be treated analytically and smoothed over for some sort of aggregate analysis. They should be brought to the forefront."

> "[Outliers] shouldn't just be treated analytically and smoothed over for some sort of aggregate analysis. They should be brought to the forefront."
>
> **JOEL K. SHAPIRO**

Get Some Context

It is important to note that Shapiro only learned of the exceptional teacher after disclosing his results to the state's board of education. This highlights another critical step in utilizing data: business leaders must work closely with data scientists to interpret these outliers, as often the leaders are the only ones with the necessary institutional knowledge or business context.

"It's unfortunate when the business person, who is relying on the analysis, gets these aggregated trends without having those outliers identified to him or her by the data scientist," Shapiro says. "The interaction with the data folks is critical so that the business experts can determine what's really going on."

Use Outliers to Optimize Automation

The Spanish-instruction example hints at yet another important lesson about outliers: one of the biggest opportunities comes in knowing which aspects of the customer experience are ripe for automation, and which are best kept in-person.

Automation can produce a more consistent customer experience—that is, fewer outliers. Depending on the circumstance, this can be a good thing or a bad thing.

“Automation tends to create a ceiling on the best experience you can have and a floor on the worst experience,” Shapiro says. “It drives everybody to the middle.”

Think of it this way: If your cell phone breaks and you need to call your service provider, would you rather wait several minutes to speak with a friendly representative or have immediate contact with a chatbot—an artificial-intelligence program that can answer basic inquiries quickly?

“Today’s chatbots may be able to solve my problem, but they probably won’t make me feel great about the experience,” says Shapiro. “If I’m talking to a live person, they might say, ‘Hey, you’re in Chicago today. My goodness, I hear it’s -3 and the wind chill is -40. Is everyone okay?’ A human touch can provide an emotional connection that most automation can’t rival.”

Outliers can be useful for helping companies determine where a given customer interaction should fall along the continuum between in-person and automated, says Shapiro. “Every organization needs to answer for itself where it wants to inject variation and where it wants to standardize or automate that variation away."

In other words, where do the benefits of an in-person interaction have the potential to outweigh the costs?

Know Your Brand

Shapiro recommends companies determine how each model of customer experience fits into its brand identity. For example, Oregon-based Dutch Bros. Coffee, known for its "bro-istas"—men and women who memorize customer preferences and dole out free coffee to down-on-their-luck patrons, leans toward personalization. Domino's, on the other hand, has reemerged as a U.S. pizza-market leader by automating customer orders through its mobile app.

"Automation, done well, can be highly efficient. Most places would say that when they fulfill orders, they're okay with standardizing or automating to be quick and accurate," Shapiro says. "But there are other elements of that customer journey that might merit a lack of automation."

It also depends on a company's desire to give up consistency in favor of a personal touch. Shapiro appreciates that many chain hotels now offer automated check-in service. But he also understands why others would want a front-desk clerk handling that job. "It could be important for them to establish a relationship," he says. "For example, me chatting with the nice Midwesterner behind the counter of the Hilton hotel where I'm staying might be a really important part of their brand."

That personal touch comes at a cost in staffing, of course. And, if those people are not trained properly and coached regularly, they can veer from established protocols or otherwise provide poor customer experiences. Companies with brands known for exceptional service might forego automation, but in their quest for that personal touch, they will need to invest in minimizing bad outcomes.

"The outliers indicate what your problems are and where your greatest successes are," Shapiro says, "and they allow you to facilitate those great successes and eliminate the problems."

Featured Faculty

Joel K. Shapiro, Clinical Associate Professor of Data Analytics at Kellogg

Learn more from Joel Shapiro in our **Executive Education programs**.

Based on insights from **Brett Gordon**

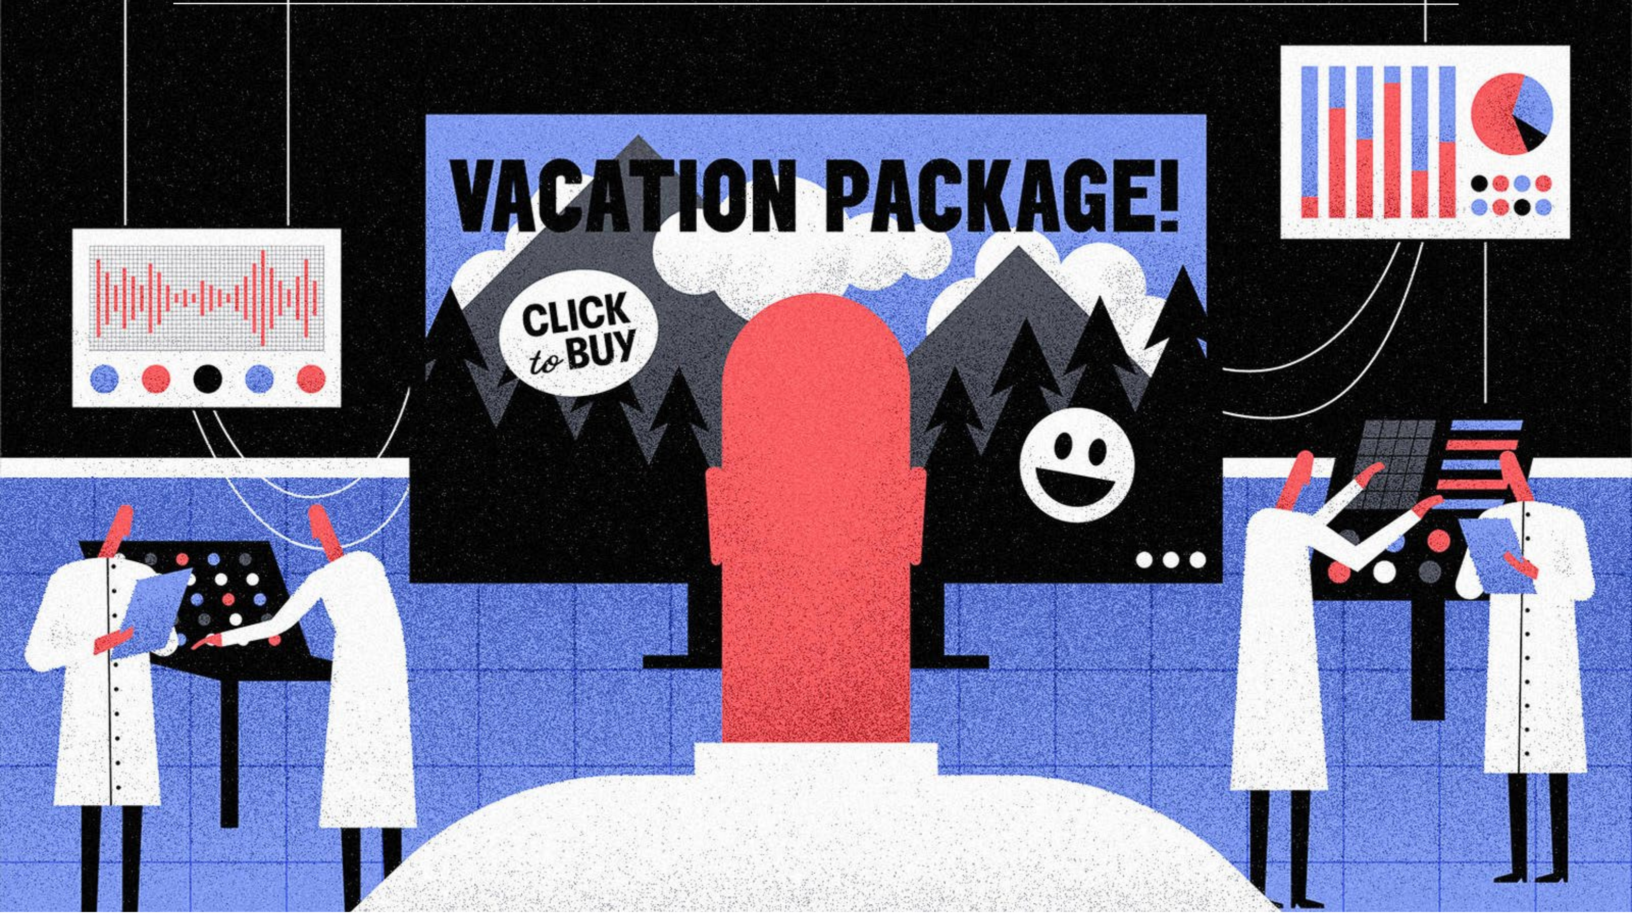

Using Experiments to Get the Most Out of Your Digital-Advertising Budget

Brands are demanding evidence that campaigns are working. Here's what they should be asking.

DIGITAL ADVERTISING CONTINUES its march toward dominance. In 2017, digital spending surpassed television advertising in the U.S. for the first time; by 2019, it overtook advertising on all traditional channels. And some estimates suggest that by 2023, digital could account for as much as two-thirds of ad spending.

This shift has, of course, impacted the industries and channels that have long relied on advertisers' dollars. But it has also drastically reshaped marketing itself, giving marketers an unparalleled amount of data upon which to make decisions—as well as a new mandate to deliver quantifiable results.

The days when a firm might simply entrust an agency with a fixed advertising budget and hope for the best are long gone, says Brett Gordon, an associate professor of marketing. "With the digital revolution, better measurement has become more feasible, and that has created more of an onus on the advertising ecosystem to demonstrate value."

So how can companies make sure they are getting the most from their digital spending? We spoke with Gordon about the importance of running randomized experiments, the analytics tools that are available to help advertisers today, and the trends that are most likely to change digital advertising in the future.

This interview was edited for length and clarity.

KELLOGG INSIGHT: As digital advertising is maturing, advertisers are increasingly demanding to know whether they are getting a return on their investment. What questions should CMOs be asking themselves to ensure they are making smart decisions?

BRETT GORDON: I think the number one thing they should always be asking is: Are they getting what they want for their money? Is it worth it? And how do they know it is worth it?

This is where the importance of correct measurement comes in. As an advertiser, you are the one who cares the most about whether the measurement is being done correctly or not. Someone selling you ads may have incentives that don't align with yours.

In many ways, measurement is very, very simple. To measure anything, you need a control group, where you are either not running the campaign or running it at a much lower volume. And you need that control group to be as similar as possible to the group where you are actually running the campaign. So if someone is telling you about the effectiveness of your campaign, you should be able to ask, "What is the basis of comparison? What was the control and why should I believe that the control is similar to the people who were in the test?" The more similar you can believe those are, the more valid the results are.

INSIGHT: That makes sense. But can you explain just how similar a control group needs to be in order to reliably measure whether your campaign is working? On the one hand, there are randomized controlled trials—experiments with true control groups, not unlike ones you would find in academic or medical research. But there also exist a plethora of alternative approaches, workarounds really, such as creating a control group (or market) that is similar in a variety of ways to the one that sees the ad, or comparing the same group before versus after a campaign.

GORDON: We have research we did with Facebook that looks at this. We were wondering if companies could rely on statistical models using observational data, i.e., not data generated from an experiment but simply data collected in the normal course of running a campaign. Our question was: How well do these observational methods do? And the beauty of partnering with Facebook was they let us directly compare randomized experiments with several observational approaches to assess whether they were a good substitute. And by and large, we found that unfortunately you couldn't rely on the observational techniques to recover anything close to the true experimental effect of the ad campaigns.

Facebook should have been a best-case scenario for these other approaches, because we had a very controlled environment and lots of information about individual users, their exposure to ads, and the purchases or conversions on or around the platform. So if we couldn't recover the actual effect using observational techniques in our nice, clean, walled garden, it would be much harder, we think, for someone else to use these observational techniques outside of this platform.

INSIGHT: How difficult is it to run these randomized controlled trials on these platforms? Is it always worth the time or money?

GORDON: These are much easier to run than they used to be. Different platforms have started to make a number of experimentation features freely available to advertisers, such as "ghost ads" testing on Google or Conversion Lift on Facebook. Companies have done a lot, I think, to make them as user friendly as possible.

But it really comes down to a different question, which is: What do you think you know versus what do you think you don't know?

As an advertiser, if you are absolutely positive that this campaign you are about to run is going to make you a lot of money, then you may not want to run an experiment. An experiment entails not showing that campaign to some people. If you are sure it is going to make you money, you want to show it to everyone possible: you don't want that control group to not see your campaign.

On the other hand, if you're uncertain, or you want to know how effective

"Brands are looking closely at their marketing budgets and asking, 'For every dollar I spend, am I getting more than a dollar back? And if that's not the case, then should I be spending it elsewhere?'"

BRETT GORDON

the campaign will be, then you need to have a control group. Only by having a control group do you actually know the counterfactual: what would have happened had you not run the campaign.

INSIGHT: OK, so you run an experiment with a randomized control group, and now you know how effective that particular campaign is. What do you do next? You don't have that same level of data about the other channels in which you may be interested in advertising, such as print or television.

GORDON: That's right. A bigger challenge is definitely knowing what to do with the information you get. I think that's where a lot of advertisers haven't cracked the nut yet.

Brands are looking closely at their marketing budgets and asking, "For every dollar I spend, am I getting more than a dollar back? And if that's not the case, then should I be spending it elsewhere?" But now they're spending it in lots of different buckets, and in some of those buckets it's easier to measure the returns than in others. So how do you, as a brand, incorporate this ambiguity into your decision-making?

This is where your own intuition and instinct come in. But by virtue of what you know about how your customers behave and how the channel behaves and who is going to see the campaign or not—data you should be able to get—you can have a sense of what might work. You will also be able to look at outcomes after the fact. Do you see a bump in your business from the right channel?

Still, now you have to weigh, say, a guaranteed 20 percent return from a digital campaign against something that you think was positive but could be more or could be less. That's why, in the end, people diversify their allocations, because they know they need to take a portfolio approach. Marketing spending is essentially a portfolio-optimization problem, except it's in some ways harder than finance, because we can't always measure the return even after the fact.

INSIGHT: How will advances in machine learning change digital advertising?

GORDON: Machine learning has already changed advertising by improving predictions about how someone will respond to an ad. Machine learning is very, very good at problems where you have lots of information and are trying to make a prediction, such as, will someone click on an ad if I show it to them. And machine learning is very good about doing this quickly and on a large scale. So it is already part of the plumbing now in a lot of digital-advertising campaigns, whether people know it or not.

INSIGHT: Are there other trends marketing leaders should keep an eye on?

GORDON: Absolutely. First is the emergence of other companies in the advertising space. Take Amazon. There's long been this duopoly of Google and Facebook, with 37 percent and 22 percent of the market, respectively. Amazon is still a long way away in terms of market share—at about 8 percent—but they have a lot of momentum. Being a retailer, they are the last touchpoint. And so as an advertiser, there is great value in advertising on their platform, because you are so close to the buying stage that the targeting capabilities become quite compelling. Or imagine if Netflix all of a sudden got into the ad business. They have fantastic customer data, and they're technologically very capable as a company. I think a lot of advertisers would love to be able to advertise on Netflix.

Second is GDPR. So far, the broader implications of this regulation have not really been felt, because implementation has been so slow and varied across websites. Many websites don't follow the letter of the law or the intent of the law. I think the near-term effects—which haven't been much—are not indicative of what the long-run effects are going to be.

Third, I think another unknown in terms

Learn more about digital advertising effectiveness from this episode of our podcast, The Insightful Leader.

of privacy is what Apple might do. Apple has lots of control over the information passed to advertisers on iPhones or in Apple apps. Companies that collect cookie data normally see all of people's movements across the web, which allows for increased targeting and optimization. But Apple may want to break up the link between people seeing an ad and people buying something. This is already part of an experimental feature in Safari, and they are in the early stages of proposing it through the W3C, which develops standards for the web.

Apple can do this because they don't currently rely on ads for a big part of their business. So in some ways, anything that they do to make it harder for the rest of the ad ecosystem to operate is probably a win for them.

Featured Faculty

Brett Gordon, Associate Professor of Marketing at Kellogg

Related Resources

Read: Is Your Digital-Advertising Campaign Working?

Quiz: How Well Do You Understand Digital Advertising?

Want to Understand Customer Sentiments in Real Time? Here's a New Tool for Analyzing Social Media Data.

Perceptions of your brand can change quickly—
but there's a novel way to keep up.

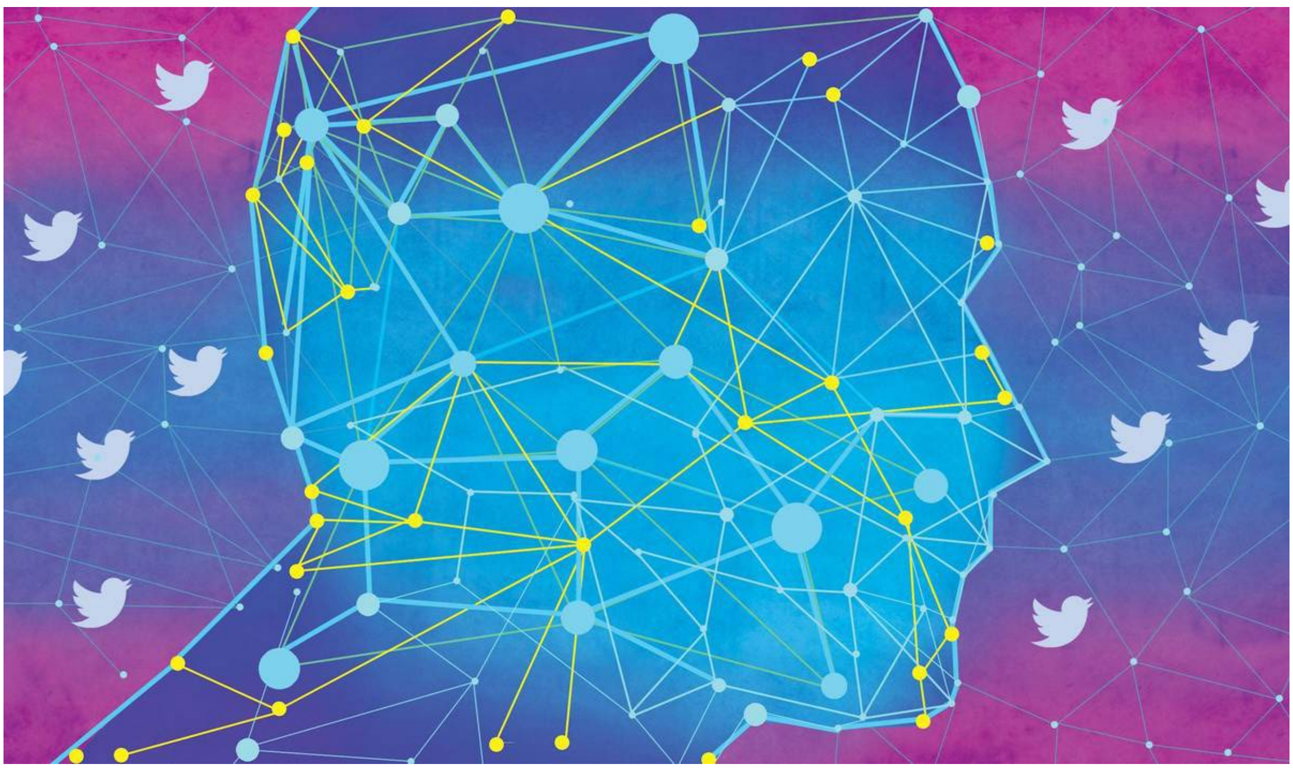

FOR DECADES, MARKETERS HAVE RELIED on surveys to gauge how customers perceive their brands. While this tried-and-true method does a good job of revealing how brands stack up against the competition on everything from health to luxury, it is also time-consuming and labor-intensive. By the time you have survey results in your hand, they may already be out of date.

Jennifer Cutler, an associate professor of marketing at the Kellogg School, thinks it may be time to send many surveys into a well-deserved retirement. Instead, she and a coauthor have developed a real-time brand-perception measurement tool based on Twitter activity.

Cutler and coauthor Aron Culotta of the Illinois Institute of Technology have created an approach that allows marketers to track in real time how their company compares to others for any attribute that interests them: in minutes, a marketer can know, for example, whether customers see Tesla as more or less luxurious than Porsche, a task that previously might have taken weeks or even months to complete. This is accomplished not by tracking what users are posting to Twitter, but rather whom they follow—an approach Cutler believes offers deeper and more nuanced insights into how companies are viewed.

It's an approach that can also be used to track emerging trends or the latest hashtags.

"There's a lot of excitement in the field of marketing about the potential to extract insights about consumers from these data, but there's definitely been a struggle to figure out how to do that," Cutler explains. Thus, much of those data remain untapped by marketers. Thanks to research like hers, however, "a lot of the barriers to entry and a lot of the obstacles to applying large-scale data mining for marketing insights are falling down."

The Power of Social-Media Data Mining

When marketers look to social media, they are often focused on what consumers are saying about their brands. Though Cutler believes text analysis has its place, there are serious drawbacks to relying on text alone. For example, although 20 percent of U.S. adults have Twitter accounts, fewer than half post actively.

"Among those who write, very few are going to write about a brand, and even fewer still are going to write about your brand," Cutler explains.

But consumers reveal a lot about themselves online, even when they say nothing at all.

These Twitter lurkers are following other users—companies, politicians, celebrities, friends—and making hand-curated lists of accounts organized by topic ("sports," "science," or "politics"). And unless they have made their Twitter account private, all of this information is publicly available.

Across these many millions of user-curated lists, certain commonalities begin to emerge. @ESPN, for instance, might appear on many user lists labeled "sports." Cutler's algorithm identifies exemplary accounts for particular topics. It searches for accounts that appear on many lists labeled, for instance, "environment" and narrows those accounts down to the strongest exemplars, such as @SierraClub or @Greenpeace.

The algorithm then looks for overlap between the followers of those exemplary accounts and the followers of a particular brand (say, Toyota Prius). This information is used to compute a score that shows how the brand is associated with the attribute. Lower scores mean most customers do not associate the brand strongly with the attribute (say, Walmart and luxury); higher scores indicate a stronger association (Toyota Prius and the environment).

When the researchers compared their computer-generated results with traditional survey results for 239 brands, they found that, in most cases, the survey results closely matched the results produced by the algorithm.

And in contrast to the sluggish process of administering surveys—or, for that matter, other automated tools that must be trained on "labeled data," or data that are painstakingly hand-tagged with information about users or their messages—the algorithm responds more quickly to shifts in public perception.

"Anytime we want to run this model, we can just query again, and if there are new players in the field—new, trendy sustainability exemplars—then we'll catch them with the new query," Cutler says.

Discovering Emerging Trends and Hashtags

The ability to identify these exemplar accounts, and quickly, can be useful for other purposes as well, like identifying trends as they emerge.

Take hashtags. Because they can pack a lot of punch into a short character count, they have become particularly important for understanding a user's perceptions and intent. As Cutler explains, a user might "give an opinion like, 'I love days like this #EARTHDAY.' And really, the hashtag is the only thing that's giving you context about what you love."

But hashtags can change seemingly overnight. Nobody wants to build a model to track engagement for the hashtag #ECOFRIDAY when there is a good chance that users will abandon that hashtag and start using

> “The real innovation here is that instead of trying to generate all of these difficult-to-obtain training data to build a model, we’re finding sources that users are already providing.”
>
> **JENNIFER CUTLER**

#SUSTAINABLESATURDAY instead before the analysis is complete.

Enter the exemplar approach. “If we can get a whole bunch of accounts that we know will be speaking about the environment at a much higher rate than average, we can then compare the language they use against a random pool or a well-defined control set of similar accounts, but which aren’t known for their environmental friendliness,” says Cutler.

Through that comparison, you can tell which hashtags or words or messages signal interest in the environment right now—as opposed to the ones that might have signaled that six months ago—and capture new signaling language as it emerges.

Cutler’s approach has other uses too. It can help marketers to determine which platforms are the most popular for certain topics of conversation, for instance. That way, you don’t waste effort on a Facebook campaign if the discussions most relevant to your brand are all on Instagram instead.

While there are currently commercial tools that track conversations and influencers, none of them really reveal how they are arriving at their results. In contrast, Cutler is very open about the algorithms she uses.

“These are algorithms that we’re publishing in journals, and so they’re pretty transparent and shockingly not very complicated. The real innovation here is that instead of trying to generate all of these difficult-to-obtain training data to build a model, we’re finding sources that users are already providing,” says Cutler.

Cutler hopes as social-media data mining becomes more accessible to marketers, it will allow them to gain deeper insights.

“As we develop these new techniques, it can start to open the door to new types of questions that marketers can ask that they haven’t been able to ask before,” she says.

Social-media data can be a boon to marketers and researchers. What happens to individuals?

Individuals badly underestimate the amount of information that can be gleaned from even the most pedestrian online behaviors. “Even really innocuous things such as ‘liking’ a supermarket’s page can eventually build a very predictive profile of a person, including basic demographics, but also more sensitive things like someone’s political leanings, religious preferences, and health conditions,” Cutler says.

Moreover, data collection is cumulative, and not always in ways that are easy to explain, track, or predict. “A consumer might not care that a specific small app can access her iPhone photos,” Cutler says. “But when she realizes that companies are merging together, or merging their data, suddenly there’s this data beast that knows everything about you.”

Read more about the connections between online behavior and individual privacy here on Kellogg Insight.

Featured Faculty

Jennifer Cutler, Associate Professor of Marketing at Kellogg

Based on research from **Artem Timoshenko**

How Companies Can Mine Online Reviews for Product-Development Gold

User-generated content contains valuable consumer insights. You just need the right technique to uncover them.

CONSUMER PRODUCTS RARELY remain static. They continually evolve based on customer needs.

Take the electric razor. About a century ago, one of the first models—the Vibro-Shave—entered the U.S. market. This razor, which was marketed to both men and women, had a vibrating handle that moved a blade from side to side. And, if you were so inclined, you could swap out the top for a massage head, to "coax" away your wrinkles.

Needless to say, razors have changed over the years.

So how does product development ordinarily occur? Companies have long relied on market research to determine how customers are using their products and whether they have underlying needs that a new feature or innovation might meet. Much of this research has traditionally involved interviews or focus groups with customers, who share how they use a product, what they like, and what they don't like. Companies then synthesize this feedback, to determine if customer needs are being met, and act on this knowledge.

But interviews and focus groups are expensive, and they can take an enormous amount of time, says Artem Timoshenko, an assistant professor of marketing at Kellogg. "Being on the market with a new razor half a year before your competitor gives you the edge."

So Artem Timoshenko and his colleague John Hauser of MIT Sloan wondered whether it was possible to glean similar insights about customer needs from *existing* customer feedback—namely, user-generated content like Amazon reviews or social-media data.

They had two specific questions: First, could professional market-research analysts extract useful information from these reviews? And second, could machine-learning algorithms enable them to do so more efficiently?

Mining Product Reviews

To address this first question, the researchers brought in a marketing-consulting company called Applied Marketing Science, Inc. (AMS). AMS has over twenty years of experience in market research and customer-need elicitation, and they had recently conducted a for-client interview-based study of customer needs for oral-care products.

"It was very convenient from both a business and a research perspective," Timoshenko explains, as toothbrushes represent a fairly standard product category and one with plentiful Amazon reviews. Moreover, AMS was excited about the researchers' questions, and the company was eager to partner.

When it comes to oral-care products, many customers report needs that are fairly straightforward: the products need to keep their teeth clean and white, keep their gums healthy, and not damage any previous dental work. But other customers might mention less expected needs, such as knowing how long to spend on various parts of their mouth during their oral-care routine. This might lead to product ideas, such as toothbrushes that beep at timed intervals or shut down after a certain number of minutes.

The experiential interviews conducted by AMS revealed 86 different customer needs for oral-care products, a typical number for such a product category. The goal of analyzing these customer needs is to find a hidden gem: a need that is very important, but that existing products do not meet well.

To determine whether marketers can glean the same kind of information about customer needs—and potential hidden gems—from user-generated online reviews as they can from interviews and focus groups, the researchers randomly selected a subset of Amazon reviews for oral-care products and provided that to a group of analysts at AMS. These analysts were not the ones who had collected or analyzed the customer interviews, but they were similarly trained. Each of the reviews in the subset was presented to the analysts in its entirety, and together the reviews added up to 12,000 sentences—which took the analysts approximately the same time to review as a standard set of 20–25 experiential-interview transcripts.

Going into the study, Timoshenko and Hauser thought that the Amazon reviews might have some advantages over traditional customer interviews. For example, perhaps they offered access to a population of customers who were unlikely to participate in a focus group.

“We could imagine that, if a company’s located in Boston, they would mostly interview Bostonians,” says Timoshenko. “But maybe people in other areas have different product experiences and usage models.”

Another possible advantage is that customers tend to write online reviews immediately after using something. Participants in a focus group, on the other hand, might have used the product a month or two before they are interviewed and have already forgotten key parts of their experience.

However, the researchers also suspected that online reviews might

> At least for some categories, we are able to fully eliminate the need to conduct interviews and focus groups. And that is the most time-consuming part of market research for customer needs.
>
> **ARTEM TIMOSHENKO**

have a major disadvantage. Specifically, "there is a lot of research suggesting that online reviews are skewed toward extremely positive or extremely negative," says Timoshenko. "So we might be missing some of the customer needs that are usually expressed in more neutral language."

For instance, the fact that a toothbrush actually cleans teeth—an important but by no means thrilling use—might not be the kind of thing that a customer would bother mentioning. That was a major concern, as articulating the entire set of customer needs can help product-management teams to identify new product opportunities, even when some of the customer needs are not surprising *ex post*.

So what did the researchers find? First, almost all—97 percent—of the customer needs identified in the interviews and focus groups were also found in the Amazon reviews.

"That immediately suggests that, at least for some categories, we are able to fully eliminate the need to conduct interviews and focus groups," says Timoshenko. "And that is the most time-consuming part of market research for customer needs."

The second finding was that the Amazon reviews contained eight additional customer needs (nearly 10 percent of the total), which were *not* mentioned during the interviews. These were not materially different from those that were mentioned by customers—they appeared to be just as important to customers and useful for future product development—suggesting that analyzing user-generated reviews could provide a more exhaustive insight into customer needs.

Timoshenko suspects that, if additional interviews and focus groups had been conducted, these needs would have eventually emerged. "But doubling the number of interviews you conduct is much more expensive, in money and time, than just doubling the amount of online content we review."

Machines Aiding Humans

Next, the researchers attempted to see whether they could use machine learning to make the human analysts more efficient. Specifically, they built an algorithm to "prescreen" the reviews, weeding out less helpful ones so that analysts could make more productive use of their time.

The researchers trained an algorithm to prescreen the reviews in two ways: it removed non-informative sentences, and it reduced redundant ones. Non-informative sentences, which make up nearly half of all of the sentences in the corpus, might simply say, "My son loves this product"—a perfectly legitimate sentiment, but not one that will lead to product innovation. Redundant reviews, also prevalent in the corpus, mention the same deficit or perk over and over again.

The researchers found that the prescreening by their algorithm allowed the analysts to find the same number of customer needs in about 20 percent fewer sentences.

“This was the proof of concept,” says Timoshenko. He is confident that with more experience and engineering, efficiencies would continue to increase, just as the methods for traditional-interview-based market research have improved over years of practice.

To that end, the researchers have made their code freely available to companies and are eager to learn about how it is being further developed and applied by companies in different industries.

One company in the food industry, for instance, has used the researchers’ methods and found that they identify very different kinds of customer needs depending on whether they search online reviews or social-media data.

Timoshenko says this highlights the fact that, as multiple sources of feedback are considered, the need for machine-learning tools will only grow.

“There is even more need for preprocessing this information,” he says, “because there are millions of Amazon reviews for a particular product—but if you want to combine that with the social-media data and online reviews from other sources, it just blows up the amount of content you have to process. And that makes machine learning very important.”

Unexpected Benefits

In doing their research, Timoshenko and Hauser found that analyzing user-generated content has another, quite unexpected, advantage over traditional interviews and focus groups: the ability to "follow up" on an intriguing customer comment or need in order to dig deeper.

In a traditional-interview setting, he explains, "you don't have the chance to call back the same interviewee and talk about this experience. It's a lost opportunity."

With user-generated content, on the other hand, you actually can explore further. With an interesting lead in mind, you might go back to the entire corpus of thousands of reviews to search for additional clues. "You don't go to exactly the same customer review, but you could look for the keyword, or a particular phrase, or the particular experience," Timoshenko says.

Overall, he wants marketers to understand that machine learning can be a powerful tool—not just for replacing human intelligence, but for augmenting it.

"One of the big breakthroughs in this research was when we agreed on the idea that machine learning cannot solve all the challenges of this process," says Timoshenko. "Most people, when they think about machine learning, they look for completely automated solutions. It appears that humans are just much better, naturally, in some tasks than machines. And they will stay better in the foreseeable future. And formulating customer needs is one of these tasks."

A customer might say, “I don’t like this toothbrush because it doesn’t have a 30-second timer.” But the underlying customer *need* is wanting to know how much time to spend on various parts of your dental routine.

“It’s very abstract. It’s very conceptual what the customer really wants. So this step is better done by humans, who can really learn and understand the human experience of other customers.”

Featured Faculty

Artem Timoshenko, Assistant Professor of Marketing at Kellogg

Related Resources

Read: People Are Tweeting about Your Products. Will It Boost Sales?

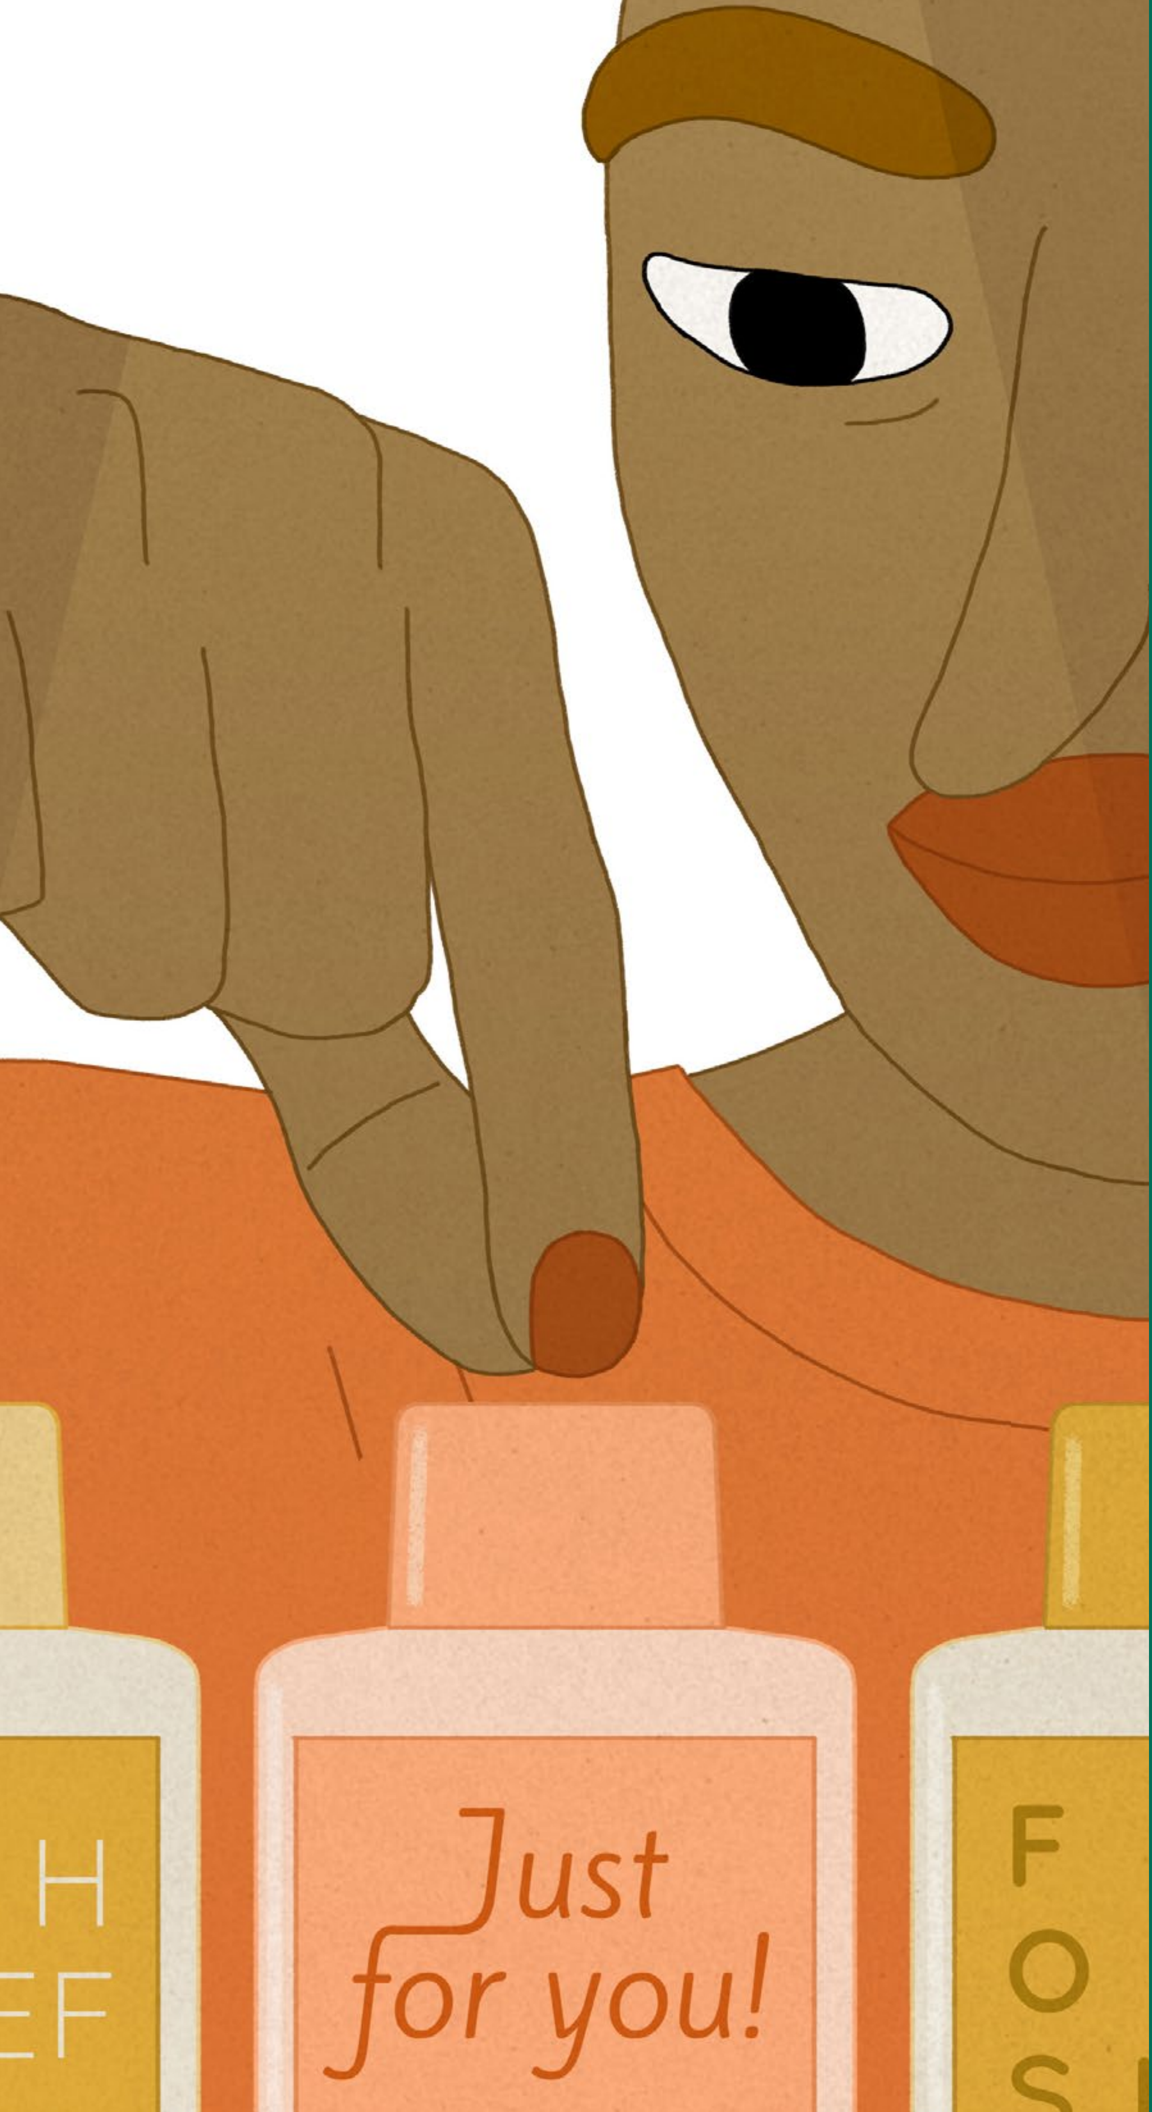

Trends

Four Ways AI Is Already Improving How Companies Serve Customers

How businesses are using these technologies to get the data they need and generate smarter recommendations.

FROM SAVVIER CHAT BOTS to automated financial planners, every day seems to bring news of yet another way artificial intelligence is changing how businesses find and serve their customers. But what's the hype, what's happening, and where is all this headed?

Adam Pah is a clinical associate professor of management and organizations at Kellogg, as well as a researcher at the Northwestern Institute on Complex Systems (NICO). Over the years, Pah has collaborated with a number of organizations as they integrate AI into their operations.

Here he highlights some of the most significant trends in AI that are changing how companies market to customers.

Using Natural Language and Video to Glean Customer Insights

People communicate mostly through natural language—not equations or Excel spreadsheets or structured daily reports. This has historically posed a challenge for marketers, because while the words we use are full of consumer insights, they have been generally inscrutable to analytical tools.

This is, of course, changing. With natural-language processing and machine-learning software, companies can now scrape vast libraries of text to generate insights on almost any subject. It has now become commonplace for companies to analyze, say, customers' Facebook comments to design more effective digital-marketing strategies.

But organizations are becoming more creative and ambitious in thinking about how to apply insights from unstructured data, according to Pah.

Consider a cell-phone manufacturer launching a new phone.

Because these manufacturers make most of their sales through service providers such as ATT or T-Mobile, they are not directly connected to their customers. This means that post-launch customer feedback has historically been pulled from either customer returns or the manual monitoring of social media and tech websites by a small group of employees.

Pah worked with a manufacturer to create a different process. "We wrote a program that crawled the Internet for mentions on Twitter and Facebook along with about 30 sites where people review and talk about cell phones," he says. "We then built a machine-learning solution that would automatically identify when the company's product was being talked about, which features of the phone were being talked about, whether they were being talked about as positive or negative attributes. Then the program would group like statements together to give an idea about the magnitude of each problem—and it did this continuously."

This new software allowed the phone manufacturer to start responding to issues immediately upon launch instead of after the typical six-week lag. It also provided a wealth of consumer insights that could be translated into future research and development.

Similar machine-learning efforts are well underway for images and video as well. There is a company, for instance, that monitors social-media feeds, televised sporting events, and shopping sites to identify when a logo appears in an image, video, or GIF. Another company analyzes and tracks the features of clothing popular with customers to determine which are trending. They then turn these insights into new styles, debuting over 1,000 new styles per week.

"You can now expand your focus beyond

Hear more from Adam Pah on this topic in this episode of The Insightful Leader podcast.

what people write and track Instagram as a way to understand when and where you are showing up, as well as whom you should be targeting," says Pah.

Making Smart Recommendations

Making recommendations is one of AI's "oldest friends," Pah says. Consider Amazon and Netflix matching customers' past behavior with other products or movies they might like.

Pah says the push to make more accurate recommendations is showing no signs of slowing—and here is another area where advances in natural-language processing and image recognition are pushing the industry forward.

Olay, for instance, has an app that allows customers to upload a photo of their own face so that its algorithm can suggest relevant skincare products. That's a key advance, he says, given that most recommendation engines are still based largely on what the consumer has previously bought. This can be a problem if customers have never before purchased a product in that category, are interested in truly novel products, or haven't been happy with their earlier purchases.

For instance, when it comes to skincare products, "I generally get whatever is on sale," Pah says. "So my past purchases are awful if you're trying to help me actually take care of my skin."

Pah points out that the art of using image recognition to generate recommendations is still being perfected: Olay's app still requires users to manually provide the app with a lot of information about their skin, for example. "We're not entirely there yet," he says.

But the idea is right. "It's actually pretty neat that they're even trying," Pah says.

Privacy Concerns Are Growing. Companies Are Responding.

Another step in providing more useful recommendations involves incorporating data from a wider range of sources than we have seen historically. "So far, we haven't moved much beyond this platform-centric world: some website recommends something else on the same website," Pah says.

There are some signs this is changing. A few years ago, for instance, Pah received an email from TurboTax asking whether he would be interested in a machine-learning product that reviewed his tax returns, credit reports, and bank-account information to provide him with customized financial-planning advice.

"That's a step in the right direction," he says. "We're starting to recognize that we need to bring in data from everywhere to make more specific and useful recommendations."

But this would ultimately require an increase in cross-platform data sharing, a prospect that many—including customers, regulators, and even the platforms themselves—are starting to reconsider.

> Typically, there is a trade-off between the usefulness of an AI tool and the amount of privacy a user enjoys.
>
> **ADAM PAH**

"Typically, there is a trade-off between the usefulness of an AI tool and the amount of privacy a user enjoys," says Pah. In the past few years, "the data have become more siloed. Some of that is in response to the European Union; a lot more of that is in response to sifting through the stones of the past and realizing that companies made a lot of errors in how they handled private information."

"Facebook still knows everything about you. Google still knows everything about you. But it would be much harder to start today, acquire the same amount of information about people as those players, and justify that collection."

In response to restrictions on data sharing, Pah sees some companies turning to creative solutions—like synthetic data. Some large retailers, for instance, are using the customer data they do have to create millions of "unbelievably realistic" synthetic customers, Pah explains. These customer profiles are then fed into other deep-learning algorithms that learn how likely they are to respond to various marketing efforts. This additional data give the retailers the scale they need to improve their algorithms.

Take a straightforward example, such as knowing which coupons to send to which customers, when, and in which format. Instead of collecting a lot of information about a given customer from a range of sources, the retailer might generate numerous synthetic customers from the bits of information they do have and play out a number of scenarios to determine which coupon is most effective.

Synthetic customers still require some real data to generate, however. While long-standing companies like Sears—which have decades of data from loyalty programs and customer credit cards—are well positioned to build composite shoppers, smaller players might struggle.

But companies that can figure it out stand to reap huge benefits. "You have unlimited and infinite customers upon which to test," Pah explains. "That, I think is, brilliant."

Understanding that Data Is the Product

The increasing value of data is also encouraging startups to consider whether their products can do double duty: fill an immediate need for customers, but also create datasets that can be licensed to more established companies.

"People are really realizing the importance of data as a product," says Pah.

Take Netradyne, a startup that uses AI to monitor whether drivers of large commercial fleets are operating safely. Its image recognition capabilities allow it to analyze dashcam video in real time, providing fleet managers with driving notifications.

That monitoring and analysis is the core product: "It can see if you're driving through a stop sign; it can see if you're driving too close to a car; it can see if you're falling asleep," explains Pah.

"The ancillary product is that they're constantly mapping the roads, so they have road-map data," he continues. These data could be useful for fleet managers aiming to make routes more efficient. Or they could help companies developing self-driving vehicles. "They're in partnership talks with a number of automotive companies."

It's only a matter of time before marketers jump on this trend—if they haven't already, says Pah.

Featured Faculty

Adam Pah, Clinical Assistant Professor of Management and Organizations at Kellogg and Research Assistant Professor at Northwestern Institute on Complex Systems (NICO)

Learn more from Adam Pah in our **Executive Education programs**.

Based on insights from **Florian Zettelmeyer** and **Inhi Cho Suh**

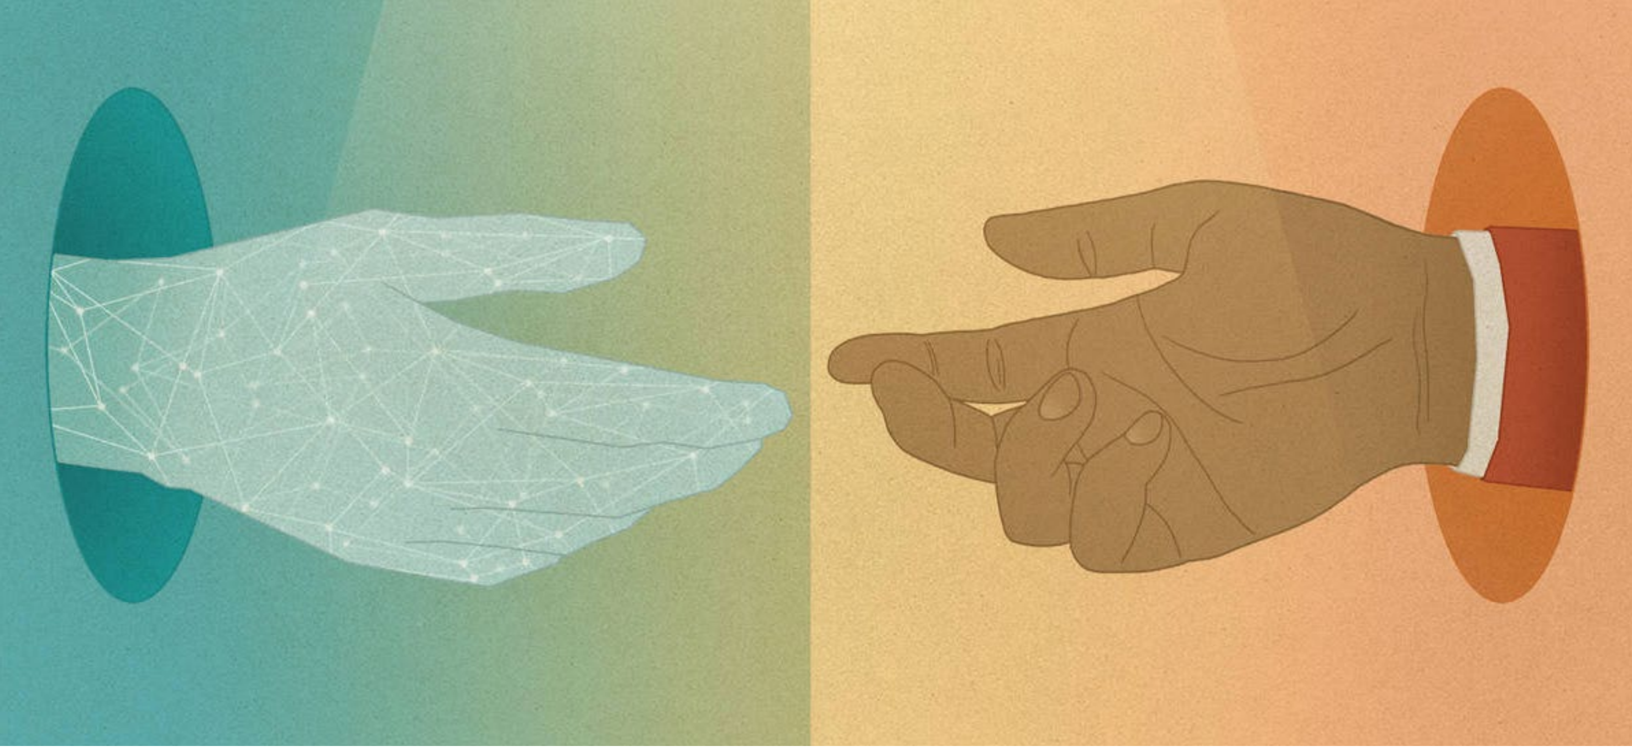

How to Build AI that Everyone Can Trust

An expert from IBM Watson discusses how to remove bias and increase transparency in machine-learning algorithms.

ARTIFICIAL INTELLIGENCE IS HERE TO STAY. Machines are getting smarter, faster, and are poised to play ever-greater roles in our healthcare, our education, our decision-making, our businesses, our news, and our governments.

Humans stand to gain from AI in a number of ways. But AI also has the potential to replicate or exacerbate long-standing biases. As machine learning has matured beyond simpler task-based algorithms, it has come to rely more heavily on deep-learning architectures that pick up on relationships that no human could see or predict. These algorithms can be extraordinarily powerful, but they are also "black boxes," where the inputs and the outputs may be visible, but how exactly the two are related is not transparent.

Given the algorithms' very complexity, bias can creep into their outputs without their designers intending it to, or without them even knowing the bias is there. So perhaps it is unsurprising that many people are wary of the power vested in machine-learning algorithms.

Inhi Cho Suh, General Manager, IBM Watson Customer Engagement, and Florian Zettelmeyer, a professor of marketing at Kellogg and chair of the school's marketing department, are both invested in understanding how deep-learning algorithms can identify, account for, and reduce bias.

The pair discuss the social and ethical challenges machine learning poses, as well as the more general question of how developers and companies can go about building AI that is transparent, fair, and socially responsible.

This interview has been edited for length and clarity.

FLORIAN ZETTELMEYER: So, let me kick it off with one example of bias in algorithms, which is the quality of face recognition. The subjects used to train the algorithm are vastly more likely to be nonminority than members of minorities. So as a result of that, the quality of facial recognition turns out to be better if you happen to look more conventionally Western than if you have some other ethnicity.

INHI CHO SUH: Yes, that's one example of a bias because of a lack of data. Another really good example of this bias is in loan approval. If you look at the financial-services sector, there are fewer women-owned businesses.

So therefore you may have loans being arbitrarily denied rather than approved because the lack of sufficient data adds too much uncertainty.

ZETTELMEYER: You don't want to approve a loan unless you have some level of certainty [in the accuracy of your algorithm], but a lack of data doesn't allow you to make your statistical inputs good enough.

What do you think of the Microsoft bot example on Twitter [where the bot quickly mirrored other users' sexist and racist language]? That's another source of bias: it seems to be a case where an algorithm gets led astray because the people it is learning from are not very nice.

SUH: There are some societal and cultural norms that are more acceptable than others. For each of us as a person, we know and we learn the difference between what is and isn't acceptable through experience. For an AI system, that's going to require a tremendous amount of thoughtful training. Otherwise, it won't pick up on the sarcasm. It'll pick up on the wrong context in the wrong situation.

ZETTELMEYER: That's right. In some sense, we face this with our children: they live in a world that is full of profanity, but we would like them to not use that language. It's very difficult. They need a set of value instructions—they can't just be picking up everything from what's around them.

SUH: Absolutely. And Western culture is very different than Eastern culture, or Middle Eastern culture. So culture must be considered, and the value code [that the algorithm is trained with] has to be intentionally designed. And you do that by bringing in policymakers, academics, designers, and researchers who understand the user's values in various contexts.

ZETTELMEYER: I think there's actually a larger point here that goes even beyond the notion of bias.

I'm trained as an economist, and very often economics has not done a particularly good job at incorporating the notion of "values" into economic analysis. There's this very strong sense of wanting to strive for efficiency, and as long as things are efficient, you can avoid considering whether the outcomes are beneficial to society.

> "In this entire space of AI and analytics, the discussion around values is supercharged."
>
> **FLORIAN ZETTELMEYER**

What I find interesting is that in this entire space of AI and analytics, the discussion around values is supercharged. I think it has to do with the fact that analytics and AI are very powerful weapons that can be used in very strategic, very targeted ways. And as a result of this, it seems absolutely crucial for an organization that chooses to implement these techniques to have a code of conduct or a set of values that governs these techniques. Right? I mean, just because you can do something doesn't mean that you actually ought to do it.

Where you have these very powerful

tools available that can really move things, you have an obligation to understand the larger impact.

SUH: Accountability is one of the five areas that we are focusing on for creating trust in AI.

Many businesses are applying AI to not just create better experiences for consumers, but to monetize for profit. They may be doing it in ways where, say, data rights may not be balanced appropriately with the return on economic value, or efficiency. So it's an important discussion: Who's accountable when there are risks in addition to benefits?

ZETTELMEYER: Do you think this is new?

SUH: I do a little bit, because in previous scenarios, business programs and applications were programmable. You had to put in the logic and rules [explicitly]. When you get into machine learning, you're not going to have direct human intervention at every step. So then, what are the design principles that you intended?

ZETTELMEYER: So a fair way of saying this is, in essence, we've always had this issue of ownership, except with machine learning, you can potentially get away with thinking you don't need it.

But you're saying that that's a fallacy, because you do need accountability at the end of the day when something blows up.

SUH: Exactly. And this goes back to [training an algorithm to have] a fundamental understanding of right and wrong in a wide range of contexts. You can't just put the chat bot into the public sphere and say, "Here, just go learn," without understanding the implications of how that system actually learns and the subsequent consequences.

ZETTELMEYER: Okay, accountability. What's your second focus area to build trust in AI?

SUH: It's a focus on values. What are the norms for a common set of core principles that you operate under? And depending on different cultural norms, whom do you bring into the process [of creating these principles]?

There's a third focus area around data rights and data privacy, mostly in terms of consumer protection—because there are companies that offer an exchange of data for a free service of some sort, and the consumer might not realize that they're actually giving permission, not just for that one instance, but for perpetuity.

ZETTELMEYER: Do you think it is realistic today to think of consumers still having some degree of ownership over their data?

SUH: I do think there's a way to solve for this. I don't think we've solved it yet, but I do think there's a possibility of enabling individuals to understand what information is being used by whom and when.

Part of that is a burden on the institutions around explainability. That's number four—being able to explain your algorithm: explain the data sets that were used, explain the approach holistically, be able to detect where you might have biases. This is why explainability and fairness—that's number five—go hand in hand.

ZETTELMEYER: In an academic context, I refer to this as transparency of execution.

I actually thought you were going to say something slightly different, that we need to move to a place where some of the more flexible algorithms like neural networks or deep learning can be interpreted.

It's a hard problem because, in some sense, precisely what makes these algorithms work so well is what makes them so hard to explain. In other words, the problem with these algorithms isn't that you can't write them down. You can always write them down. The problem is that it's very difficult to create some easily understandable association between inputs and outputs, because everything depends on everything else.

But I think the point you were making is: okay, even if we do have a so-called "black box" algorithm, a lot of the biases arise, not necessarily from the algorithm per se, but from the fact that we're applying this algorithm to a particular setting and data set, yet it's just not clear to people how it's being implemented.

SUH: That's right.

When and for what purpose are we actually applying AI? What are the major sources of that data? And how are we working to, if not eliminate bias, maybe mitigate it?

ZETTELMEYER: I think a lot of the trust problems that have occurred in the tech industry—and particularly in advertising—over the last years are directly related to a lack of transparency of that type. I'm always amazed that when you go to the big advertising platforms, and you approach them purely as a consumer, and then you approach them as a client, it feels like you're dealing with two different universes. As consumer, I'm not sure you have the same sense of exactly what's happening behind the scenes as you do if you happen to be an advertiser, and you have exposure to all the digital tools that you can use for targeting.

I think transparency, the way you're talking about it, is not particularly well implemented in many tech companies.

SUH: No. And there's not a common language for talking about it either, in terms of explicitly saying, "We only use data that we have access and rights to, and this is how we collect it, and you've given us permission for it." The standards around the language itself are still being developed.

ZETTELMEYER: What are you doing about all this at IBM?

SUH: We actually developed a 360 degrees fairness kit as part of our broader AI OpenScale initiative. AI OpenScale is an open-technology platform that enables your business with visibility, control, and the ability to improve AI deployments, helps explain AI outcomes, and scales AI usage with automated neural-network design and deployment, all within a unified management console. It includes open-source toolkits to check for unwanted biases in data sets and machine-learning modules. It checks for biases like explainability around your data sets to provide feedback on different aspects of your models.

It's the first open-platform and open-source toolkit to even begin to get developers thinking about bias proactively.

Featured Faculty and Experts

Florian Zettelmeyer, Nancy L. Ertle Professor of Marketing at Kellogg

Inhi Cho Suh, General Manager, IBM Watson Customer Engagement

Related Resources

Read: Watson Is Just the Beginning

Featured Faculty and Experts

Eric T. Anderson is the Hartmarx professor and former chair of the marketing department at the Kellogg School of Management, as well as director of the Center for Global Marketing Practice. His research interests include analytics, retailing, pricing strategy, innovation, new products, and channel management. He is cowriting a book with Florian Zettelmeyer on analytics and AI for marketing leaders, to be published by McGraw-Hill.

Jennifer Cutler is an associate professor of marketing at Kellogg. In her research, she draws on both artificial intelligence and psychology to gain insights about consumers and organizations through the analysis of social-media data. Before becoming an academic, she designed speech-recognition and natural-language-processing technologies at Microsoft.

Steve Franconeri is a professor of psychology at Northwestern University, a professor of management and organizations at Kellogg (by courtesy), and director of the Northwestern Cognitive Science Program. He leads a team of researchers who explore how leveraging the visual system can help people think, remember, and communicate more efficiently.

Brett Gordon is an associate professor of marketing at Kellogg. His research interests include pricing, advertising, promotions, retailing, innovation, and competitive strategy.

Kris Hammond is the Bill and Cathy Osborn professor of computer science at Northwestern University's McCormick School. He researches human–machine interaction, context-driven information systems, and artificial intelligence. He is also chief scientist at Narrative Science, a company focused on automatically generating narratives from data.

Eric Leininger is a clinical professor of executive education at Kellogg, where he also coleads the school's Chief Marketing Officer Program. Before coming to Kellogg, he was corporate senior vice president at McDonald's Corporation.

Tom O'Toole is a clinical professor of marketing and executive director of the Program for Data Analytics at Kellogg. Before coming to Kellogg, he had several roles at United Airlines, including chief marketing officer, senior vice president, and president, MileagePlus.

Adam Pah is a clinical assistant professor of management and organizations at Kellogg, as well as a research assistant professor and associate director for NICO (Northwestern Institute on Complex Systems). His research uses statistical and network-based models to understand how social phenomena evolve over time.

Mohanbir Sawhney is the McCormick Foundation Chair of Technology and a clinical professor of marketing at Kellogg. He is a globally recognized scholar, teacher, consultant, and speaker in business innovation, modern marketing, and enterprise analytics.

Joel Shapiro is a clinical associate professor of data analytics and academic director of executive education at Kellogg, with expertise in building organizational data-analytics capabilities. Prior to joining Kellogg, he served as associate dean of academics for Northwestern University's School of Professional Studies.

Inhi Cho Suh is general manager, IBM Watson Customer Engagement.

Artem Timoshenko is an assistant professor of marketing at Kellogg. He develops and applies quantitative methods to improve marketing practice and product development.

Jing Wang is head of the Center for Analytics and Insights at Vanguard.

Florian Zettelmeyer is the Nancy L. Ertle professor of marketing at Kellogg and the current department chair. He also founded and directs the Program on Data Analytics at Kellogg. His research addresses how the information consumers have about firms and the information firms have about consumers affect firm behavior. He is cowriting a book with Eric Anderson on analytics and AI for marketing leaders, to be published by McGraw-Hill.

Northwestern | Kellogg
School of Management