ARTIFICIAL INTELLIGENCE

The Human-AI Partnership for Better Customer Experience

Companies must carefully think through their AI strategy and decide what their goals and expectations are for both AI *and* human workers if they're to optimize customer experience.

CONTENTS

SPECIAL COLLECTION

The Human-AI Partnership for Better Customer Experience

The Future of Customer Service Is AI-Human Collaboration

P.V. Kannan and Josh Bernoff

Successful AI-powered customer service systems will depend on bots working with humans, not replacing them.

Customer service is traditionally considered a cost center, so many organizations have focused their customer improvement efforts on reducing costs. This proves to be a critical mistake, as everyone is left unhappy. Even as customers are sick of pressing two for reservations and three for service, service reps are sick of answering the same questions over and over.

The latest technology for service is virtual agents: Automated systems, trained on service transcripts, that can use AI to recognize and respond to customer requests whether by phone or chat. Given the history of customer service, you might think the focus here is once again on cutting jobs and saving money. But it's not.

Our combined experience in customer strategy and running a company that helps build virtual agent systems has demonstrated two fundamental — and counterintuitive — facts about customer service and automation.

First, the most significant gains from virtual customer service agents are from improvements in customer experience, *not* cost savings.

And second, successful virtual agent systems depend on bots working with humans, not replacing them.

Can automation actually improve service? When customers think of automation in customer service, they remember what happened with interactive voice response — a telephony system that made them navigate phone menus for gathering information but often did not provide shorter wait times or better service. Online chat with agents can be similarly frustrating, because the typical agent is juggling two to six sessions at a time, causing slow responses.

But virtual agents signal a shift for customer experience. If

a virtual agent can interpret the intent behind your chat or phone request, it can get you an answer more quickly and efficiently than a human agent. For most common queries, this delivers a better experience than interacting with an actual human.

For example, at SiriusXM Satellite Radio, a virtual agent with infinite patience will walk you through the process of initializing your new car radio. It will look up your account based on your phone number, wait while you make sure the car has a clear view to the sky, and connect to the system that sends codes down to your radio to allow it to start working. The virtual agent doesn't care if that takes three minutes or 25 minutes because it is designed to act efficiently on your intent, not to minimize the agent's time on the phone.

But not every customer service call can be easily answered by a bot alone. They succeed best because they're part of a system that includes human agents as well, capitalizing on the strengths of each.

How customer service agents and bots work together. Humans and bots have different skills. Human customer service agents easily recognize when someone is frustrated and can respond with empathy. AI-powered virtual agents, on the other hand, are wizards at assembling data from disparate systems to render a judgment instantly, even if they lack the emotional intelligence to know why such a decision might be right.

That's why companies are now designing service flows in which virtual agents and people work together. Such workflows help free human representatives from the drudgery of rooting through computer systems to find information. The human agents can instead concentrate on the exceptions — customer problems the system has not encountered before, or frustrated customers that demand empathy. The virtual agents solve the common problems efficiently, while the human agents deal with the high-stakes issues that don't fit the usual pattern.

In customer service, there are three basic modes in which virtual and human agents work together: bots initiating the conversation and handing off to a person; human agents serving customers with help from bots; or bots serving customers with human agents supervising. We'll show examples of each below.

Bots can save human agents time. Based on our experience with deploying virtual agent systems, those systems can typically handle 80% of incoming questions without assistance. But what happens when the question is too complex or the customer too upset to deal with a virtual agent?

Even then, virtual agents can save time. Some virtual agent systems are actually designed to collect information, such as the customer's name or account number and a description of the problem, and suggest resolutions as they hand off the call to a human representative. For example, at the rental car agency Avis/Budget, virtual agents collect customers' names and the places and times for pickup and drop-off, cutting 30 seconds off the call time for the human agent and eliminating repeated questions for the customer.

Handoffs between bots and people can improve efficiency. For example, at the marketing service company HubSpot, a chatbot qualifies leads, delivers content, and connects potential customers with HubSpot's (human) sales staff — all through Facebook Messenger. Having been prequalified, such leads are 40% more likely to be willing to engage with salespeople.

Bots can make human agents smarter. Some companies — like high-end professional services firms — would rather have human agents handle all service calls. But those agents can appear smarter if a bot is whispering in their ears. For example, when there's a new discount or an out-of-stock product, the bot will bring it to the attention of the agent. It's like having an assistant scurrying around to figure out the best things for you to say to the customer.

ABIE, a customer service bot at Allstate Business Insurance, was designed to help salespeople. Allstate's independent insurance agents use it to look up specific pieces of information, like what it would cost to insure a landscaping business and or what rules apply to medical labs. In the past, finding that information would require searching through hundreds of documents, but ABIE quickly brings the correct knowledge to light to make the sales staff more efficient. This allows ABIE to handle more than 25,000 inquiries every month.

Bots can improve with human supervision. In cases where virtual agents can't figure things out, a human supervisor can be a great help. When it comes to customer service through online chat, a human agent can efficiently manage eight or 10 conversations between chatbots and customers — a far higher level of productivity than if the agent were answering the questions herself. And when a virtual agent gets stuck, the more experienced human can step in and solve the problem. Not only that, but the human agent can tag the problem when the virtual agent bogs down, which allows the AI system to learn from that scenario and become smarter for the next question.

In this mode of interaction, both the bots and humans are able to focus on what they do best. For the virtual agents, that is handling routine cases quickly and efficiently. For the human supervisors, that is solving more complex problems, using empathy and emotional intelligence, and training the virtual agents on cases they didn't recognize at first.

Bots can free up people to work on more interesting challenges. At InterContinental Hotels Group (IHG), virtual agents are working to solve technical problems for employees who call the IT help desk. With a workforce of more than 30,000 employees, IHG recognizes significant gains in employee satisfaction.

IHG used machine learning to review chat transcripts from customers contacting the help desk and constructed a virtual agent that could answer common questions. The result was a system that worked *with* the 150 help desk staff instead of replacing them.

IHG's help desk can now scale to meet spikes in demand that might result from putting a new technology system in place. As Scot Whigham, who designed the system for IHG, explained, "We can scale up to meet what comes into our systems, but then scale down so that we don't have excess capacity that we are paying for." Instead of paying extra staff in times of low demand, the organization can redeploy resources to perform other important things on the IT to-do list.

For the tasks that it was trained for, like password resets, the virtual agent system can now handle 80% to 85% of the volume of questions coming in. Even when the virtual agents can't answer a question, they can help. The virtual agent system is able to collect basic information from the worker, regardless of whether it's answering the question itself or handing it off to a help desk staffer. Whigham reports that among hotel workers who have engaged with the system, satisfaction rates remain high.

The future of customer service is human-machine collaboration. As all of these examples demonstrate, the AI-driven transformation of customer service is not about getting rid of workers — it's about making them smarter. When machines handle routine inquiries, customers are happier. And when service staff can concentrate on more complex questions — or on answering questions with a bot making suggestions — they can deliver far better service.

People doing knowledge-based work have always been more productive and thoughtful with the help of computers and information. Customer service is no exception. Now customer service workers can leave the drudge work to the machines and concentrate on being the sympathetic ear and clever problem solver for the customers who really need them.

About The Authors

P.V. Kannan is the cofounder and CEO of the San Jose, California-based [247].ai, a vendor of AI-driven customer experience software and services. His new book is *The Age of Intent: Using Artificial Intelligence to Deliver a Superior Customer Experience*. His views were featured in Thomas L. Friedman's *The World Is Flat*. Josh Bernoff (@jbernoff) is the author or coauthor of six business strategy books, including *Writing Without Bullshit: Boost Your Career by Saying What You Mean*.

Four Challenges to Overcome for AI-Driven Customer Experience

P.V. Kannan and Josh Bernoff

To launch a successful virtual agent project for your organization's customer service capability, the company must overcome four hurdles.

Are you ready to turn your company's customer service over to AI-powered virtual agents? Whether your goals include creating a fully digital business, improving customer experience, or cutting costs, virtual agents and automation offer many benefits. But, as with any technology, making the switch isn't trivial.

Consider the experience of Mark Baylis, vice president of customer service and digital customer engagement at Optus, Australia's second-largest telecom operator. He believed in the potential advantages of using virtual agents for customer service, but he needed to convince the company's leadership team that the investment would make sense.

Baylis planned to have virtual chat agents handle half the incoming queries, which would cut costs by 50%. But the cost savings alone weren't enough to overcome corporate resistance. Baylis recognized that this shift was not just about the customer service division but about the company's larger digital transformation. By positioning the switch as part of the larger transformation effort, the investment in virtual agents became more persuasive to management. It also helped that virtual agents could handle more customer requests at peak hours and improve the quality and consistency of Optus' customer experience.

By focusing on strategy, not just cost savings, Baylis was able to win over Optus' CEO and C-level decision makers at a meeting in 2017. One person at that meeting described it this way: "The general reaction was excitement. It was in line with trends in the market." Leaders at the meeting recognized the need to adopt the new technology and master it to prepare for the future.

In many ways, what happened at Optus is typical for companies that are considering AI-enabled technologies for their customer-experience strategy. Proponents must

overcome four types of hurdles: economic, technical, political, and cultural.

The economic hurdle: demonstrating improvements to the bottom line. If your company currently has more than 1,000 phone or text chat agents answering customer questions, then it's a candidate to benefit economically from a virtual agent system. Because they scale without the need to hire more people, virtual agents can pay for themselves — provided there are enough customer service intents that are easy to automate.

For example, at Dish Network, where customers initiate 6 million chats per year for sales and support, virtual agent automation generated millions of dollars in annual savings. And at cable operator Charter Communications, AI-based customer service enabled the company to automate 50% of customer service chats.

You can also make the economic case by accelerating sales as you reduce costs. At Dish Network, a surge in demand for pay-per-view fights was challenging sales staff; switching to virtual agents enabled the company to recognize more than $6 million in pay-per-view fight revenue.

Consider economic benefits in the broadest possible perspective. We've seen companies justify the decision to implement AI by citing not just cost cutting but sales increases, customer service speed improvements that improve retention, and even delivering on promises to shareholders.

The technical hurdle: ensuring that systems are ready. Virtual agents need to answer questions like "What's my balance?" or "What's the status of my order?" That requires access to corporate systems, which only happens when IT agrees on priorities with the lines of the business directly implementing virtual agents.

Sometimes, technical obstacles are insurmountable. At one large online retailer, a "decisioning engine" system was responsible for resolving customers' queries. When the company tried to hook up a virtual agent system to it, over 95% of the queries returned a response of "transfer to agent." Because no single department would take responsibility for the decisioning system, the problem was unfixable. The virtual agent system never launched because the rest of the corporate technology wasn't equipped to support it.

For a customer service group to avoid such problems, it must first review the breadth of queries coming in, classifying them by *intent* — the types of problems that the customer is seeking to fix. Success is most likely if a majority of the queries come from a common set of intents, such as "check balance," "return product," "reset password," or "increase credit limit." If the queries are all over the map, automating them is less likely to succeed.

A company is a good candidate for service automation if it has existing APIs (application programming interfaces) or can build new ones that return necessary data from corporate systems with subsecond response times. Those systems need to be extremely reliable and capable of scaling up in times of peak demand. IT staff members often already share these goals for corporate system improvements. Virtual agents need such access because when it comes to delays in returning information, customers are typically less patient with automated agents than they are with human agents.

The political hurdle: winning over management. Virtual agents can affect multiple parts of the company, including its brand. As a result, the decision to commit to virtual agents typically involves not just the head of customer service, but an array of C-level decision makers, including the chief operating officer, the CIO, some of the CIO's direct reports, the head of sales, and often the chief experience officer (assuming the company has one). Because the virtual agent can become a highly visible part of the brand, chief marking officers and heads of product management may also become involved. At Bank of America, for example, the company's virtual agent Erica appears in ads and marketing emails because it reflects the company's responsive and technologically advanced culture.

The best way to corral a group of decision makers who are this interconnected and influential is to frame the project within an existing strategic initiative at the company. At Optus, as at many companies, this was the corporate push for digital transformation. Digitizing and automating customer service creates a desirable result: It brings a highly human-centric and sometimes inconsistent part of the

business under the digital umbrella.

Digital transformation isn't the only justification. At one travel company, virtual agents got the green light because the company had promised cost savings in the wake of a merger. At another, virtual agents allowed an IT department to handle more queries without expanding beyond its physical space constraints.

The cultural hurdle: getting past human-centric prejudices. In the end, much of the resistance to virtual agents comes down to a simple question: Is the company willing to let machines talk to customers, or is it a job that only people can handle? That's a cultural and emotional challenge, not a technical one.

The surprising reality is that many customers would *prefer* to talk to a computer. Just as consumers became comfortable with websites and apps, they are likely to become happier with Alexa-like responses from machines — so long as those machines can understand the problems and supply the right answers.

As we wrote in our last article, managers can become more comfortable when virtual agents and human agents work side by side to solve customers' problems. This can help enable policy shifts that lead to cultural shifts. For example, at one credit card company, policy demanded that customers had to talk to a person to file a claim for alleging an unauthorized purchase. At SiriusXM Satellite Radio, rules specified that customers couldn't cancel service without talking to a person. In both cases, the companies revised their rules once they realized that virtual agents could be as efficient and secure as human agents.

IT's not-invented-here attitudes can be another cultural obstacle. Building virtual agents from scratch requires not just AI mastery but knowledge of the most common queries in the company's vertical (for example, travel, banking, or retail). The large vendors in this space already have that knowledge and have implemented similar systems. Quick and successful pilot projects built in collaboration with knowledgeable vendors can help reduce this cultural hurdle.

How to succeed with virtual agents. Examining these four types of hurdles reveals what it will take to make virtual agents successful. Virtual agent proponents must demonstrate both strategic fit with the company and the readiness of technical systems. They must tie virtual agents to corporate priorities like digital transformation. And they must start by rolling out pilot systems that demonstrate both customer acceptance and efficiency advantages.

In the end, companies will figure out that they need to get started in implementing automated conversational interfaces to their customer service. Because, like the web and mobile apps before them, conversational interfaces like Amazon's Alexa and Apple's Siri will eventually raise customers' expectations of every service interaction with every company.

About The Authors

P.V. Kannan is the cofounder and CEO of the San Jose, California-based [24]7.ai, a vendor of AI-driven customer experience software and services. His new book is *The Age of Intent: Using Artificial Intelligence to Deliver a Superior Customer Experience*. His views were featured in Thomas L. Friedman's *The World Is Flat*. Josh Bernoff (@jbernoff) is the author or coauthor of six business strategy books, including *Writing Without Bullshit: Boost Your Career by Saying What You Mean*.

Rethink AI Objectives

Sam Ransbotham

Using AI to create humanlike computers is a shortsighted goal.

When customers feel like no one is listening, they may be right. Or, they may be unknowingly talking to a machine. Last week, Google announced Duplex, an artificial intelligence (AI) assistant that can handle customer service requests, such as booking an appointment or providing basic information. We've had automated phone attendants for years, but the new buzz is that customers can't tell that this automated attendant isn't human.

The Duplex announcement quickly brought questions about transparency. Should companies include a notice that "objects on the phone are less human than they appear?"

Of course they should (and Duplex will). Investing massively in technology to be intentionally duplicitous seems like it belongs in the junkyard of bad ideas we shouldn't even consider, such as earthquake generators or encryption back doors. What would be the benefits of being dishonest with customers? Is the "gee whiz" worth the potential backlash?

Just because we can do something doesn't mean we should. Consider the benefits of being up-front about interactions with bots:

- **Customer satisfaction can increase.** For example, if customers know they are interacting with a computer and not using another person's time, they may spend more time customizing their selection. I've certainly spent far more time online investigating travel options than I would ever subject a human travel agent to. With better customization, customer satisfaction can increase.

- **Customers can be better informed.** When interacting with other humans, people hesitate to ask what they perceive to be dumb or repetitive questions. We can be reluctant to admit that we still don't understand. But when interacting with machines, we are free to clarify to our heart's content. As the use of technology removes the concern about long call times, we don't have to worry about using a machine's time.

- **Customer interaction can be faster.** Social norms dictate that we first lubricate the conversation with bromides and banalities about weather and pleasantries. I deeply suspect that most customer service people actually don't care if "I'm doing OK for a Monday." Knowing that we are interacting with a machine, we can get straight to the point, saving everyone time.

- **Customers may be more honest.** We care what other people think about us. As a result, we may not be completely honest when what we think differs from what

we believe the other person wants to hear — the social desirability response bias. We are less likely to worry about what a machine thinks about us than what another human thinks.

So why are companies even considering the duplicitous option? Like dogs who don't know what to do once they've caught the car they've been barking at and chasing, we've chased the Turing test (which tests whether a machine can mimic a human's intelligent behavior) and are not sure what to do now that we've caught it.

For science, passing the Turing test makes sense as a goal. For business, other goals most likely make a lot more sense. Even if computers can pass for humans, human mimicry may be a shortsighted goal. We already have humans. We don't need to imbue these bots with the "hmms," "ums," and "uhs" that are part of normal imperfect speech patterns. Duplicating misses an opportunity for more.

If not "humanlike," what should the objectives be? In this context, they might be:

- Using AI to learn from prior customer calls and avoid them entirely. Encoding and decoding requests back and forth into speech may show impressive technological prowess but probably means the customer hasn't found the information they need from other channels. The AI attendant may be superfluous, automating processes that shouldn't exist at all. Customer satisfaction may increase with a clearer online appointment system that avoids the call in the first place. Instead, ongoing analysis of call logs using AI can be part of overall process improvement.

- Using AI to assist humans during a call. By analyzing call content, businesses can use that information to improve security and improve service. These are aspects better done by machines than humans.

- Using AI to anticipate calls and plan staffing. Pretending to be human is far from the first deception in customer service calls. Instead, we now endlessly hear "your call is very important to us" as we wait while the company prioritizes staff efficiency over customer time. While this may be perfectly rational, AI can better anticipate call volumes and improve both company and customer efficiency.

AI can be a phenomenal tool to improve business and customer service (specifically). But there are much better ways to use this technology than mimicking humans and fooling customers. After all, we're still supposed to be the intelligent ones here.

About The Author

Sam Ransbotham is an associate professor of information systems at the Carroll School of Management at Boston College and the *MIT Sloan Management Review* guest editor for the Artificial Intelligence and Business Strategy Big Idea Initiative. He can be reached at sam.ransbotham@bc.edu and on Twitter @ransbotham.