



How artificial intelligence will change the future of marketing

Thomas Davenport¹ · Abhijit Guha² · Dhruv Grewal³ · Timna Bressgott⁴

Published online: 10 October 2019
© The Author(s) 2019

Abstract

In the future, artificial intelligence (AI) is likely to substantially change both marketing strategies and customer behaviors. Building from not only extant research but also extensive interactions with practice, the authors propose a multidimensional framework for understanding the impact of AI involving intelligence levels, task types, and whether AI is embedded in a robot. Prior research typically addresses a subset of these dimensions; this paper integrates all three into a single framework. Next, the authors propose a research agenda that addresses not only how marketing strategies and customer behaviors will change in the future, but also highlights important policy questions relating to privacy, bias and ethics. Finally, the authors suggest AI will be more effective if it augments (rather than replaces) human managers.

Keywords Artificial intelligence · Marketing strategy · Robots · Privacy · Bias · Ethics

AI is going to make our lives better in the future.
—Mark Zuckerberg, CEO, Facebook

Introduction

In the future, artificial intelligence (AI) appears likely to influence marketing strategies, including business models, sales processes, and customer service options, as well as customer behaviors. These impending transformations might be best understood using three illustrative cases from diverse industries (see Table 1). First, in the transportation industry, driverless, AI-enabled cars may be just around the corner, promising to alter both business models and customer behavior. Taxi and

ride-sharing businesses must evolve to avoid being marginalized by AI-enabled transportation models; demand for automobile insurance (from individual customers) and breathalyzers (fewer people will drive, especially after drinking) will likely diminish, whereas demand for security systems that protect cars from being hacked will increase (Hayes 2015). Driverless vehicles could also impact the attractiveness of real estate, because (1) driverless cars can move at faster speeds, and so commute times will reduce, and (2) commute times will be more productive for passengers, who can safely work while being driven to their destination. As such, far flung suburbs may become more attractive, vis-à-vis the case today.

Second, AI will affect sales processes in various industries. Most salespeople still rely on a telephone call (or equivalent) as a critical part of the sales process. In the future, salespeople

Thomas Davenport, Abhijit Guha, Dhruv Grewal and Timna Bressgott contributed to the writing of the paper. Mark Houston served as accepting Editor for this article.

✉ Timna Bressgott
t.bressgott@maastrichtuniversity.nl

Thomas Davenport
tdavenport@babson.edu

Abhijit Guha
abhijit.guha@moore.sc.edu

Dhruv Grewal
dgrewal@babson.edu

¹ Department of Technology, Operations, and Information Management, Babson College, Babson Park, MA 02457, USA

² Department of Marketing, Darla Moore School of Business, University of South Carolina, Columbia, SC 29208, USA

³ Department of Marketing, Babson College, Babson Park, MA 02457, USA

⁴ Department of Marketing and Supply Chain Management, Maastricht University, Tongersestraat 53, 6211, LM Maastricht, The Netherlands

will be assisted by an AI agent that monitors teleconversations in real time. For example, using advanced voice analysis capabilities, an AI agent might be able to infer from a customer's tone that an unmentioned issue remains a problem and provide real-time feedback to guide the (human) salesperson's next approach. In this sense, AI could augment salespersons' capabilities, but it also might trigger unintended negative consequences, especially if customers feel uncomfortable about AI monitoring conversations. Also, in the future, firms may primarily use AI bots,¹ which—in some cases—function as well as human salespeople, to make initial contact with sales prospects. But the danger remains that if customers discover that they are interacting with a bot, they may become uncomfortable, triggering negative consequences.

Third, the business model currently used by online retailers generally requires customers to place orders, after which the online retailer ships the products (the shopping-then-shipping model—Agrawal et al. 2018; Gans et al. 2017). With AI, online retailers may be able to predict what customers will want; assuming that these predictions achieve high accuracy, retailers might transition to a shipping-then-shopping business model. That is, retailers will use AI to identify customers' preferences and ship items to customers without a formal order, with customers having the option to return what they do not need (Agrawal et al. 2018; Gans et al. 2017). This shift would transform retailers' marketing strategies, business models, and customer behaviors (e.g., information search). Businesses like Birchbox, Stitch Fix and Trendy Butler already use AI to try to predict what their customers want, with varying levels of success.

The three use cases (above) illustrate why so many academics and practitioners anticipate that AI will change the face of marketing strategies and customers' behaviors. In fact, a survey by Salesforce shows that AI will be the technology most adopted by marketers in the coming years (Columbus 2019). The necessary factors to allow AI to deliver on its promises may be in place already; it has been stated that “this very moment is the great inflection point of history” (Reese 2018, p. 38). Yet this argument can be challenged. First, the technological capability required to execute the preceding examples remains inadequate. By way of an exemplar, self-driving cars are not ready for deployment (Lowy 2016), as—amongst other things—currently self-driving cars cannot handle bad weather conditions. Predictive analytics also need to improve substantially before retailers can adopt shipping-then-shopping practices that avoid substantial product returns and the associated negative affect. Putting all this together, it appears that marketing managers and researchers need

insights about not only the ultimate promise of AI, but also the pathway and timelines along which AI is likely to develop. This paper addresses the issues above, building not only from a review of literature across marketing (and more generally, business), psychology, sociology, computer science, and robotics, but also from extensive interactions with practitioners.

Second, the preceding examples highlight mostly positive consequences of AI, without detailing the widespread, justifiable concerns associated with their use. Technologists such as Elon Musk believe that AI is “dangerous” (Metz 2018). AI might not deliver on all its promises, due to the challenges it introduces related to data privacy, algorithmic biases, and ethics (Larson 2019).

We argue that the marketing discipline should take a lead role in addressing these questions, because arguably it has the most to gain from AI. In an analysis of more than 400 AI use cases, across 19 industries and 9 business functions, McKinsey & Co. indicates that the greatest potential value of AI pertains to domains related to marketing and sales (Chui et al. 2018), through impacts on marketing activities such as next-best offers to customers (Davenport et al. 2011), programmatic buying of digital ads (Parekh 2018), and predictive lead scoring (Harding 2017). The impact of AI varies by industry; the impact of AI on marketing is highest in industries such as consumer packaged goods, retail, banking, and travel. These industries inherently involve frequent contact with large numbers of customers, and produce vast amounts of customer transaction data and customer attribute data. Further, information from external sources, such as social media or reports by data brokers, can augment these data. Thereafter, AI can be leveraged to analyze such data and deliver personalized recommendations (relating to next product to buy, optimal price etc.) in real time (Mehta et al. 2018).

Yet marketing literature related to AI is relatively sparse, prompting this effort to propose a framework that describes both where AI stands today and how it is likely to evolve. Marketers plan to use AI in areas like segmentation and analytics (related to marketing strategy) and messaging, personalization and predictive behaviors (linked to customer behaviors) (Columbus 2019). Thus, we also propose an agenda for future research, in which we delineate how AI may affect marketing strategies and customer behaviors. In so doing, we respond to mounting calls that AI be studied not only by those in computer science, but also studied by those who can integrate and incorporate insights from psychology, economics and other social sciences (Rahwan et al. 2019; also see Burrows 2019).

Introduction to artificial intelligence

Researchers propose that AI “refers to programs, algorithms, systems and machines that demonstrate intelligence” (Shankar

¹ Miller (2016) outlines the difference between an AI bot and a chatbot. In brief, chatbots rely on (relatively) simple algorithms, whereas AI bots have greater capabilities, incorporating complex algorithms and NLP.

Table 1 Select use cases (in the order in which they appear in the paper)

Industry or Usage Context (specific firm or AI application)	Description
AI in driverless cars (e.g., Tesla)	In the future, AI-enabled cars may allow for car journeys without any driver input, with the potential to significantly impact various industries (e.g., insurance, taxi services) and customer behaviors (e.g., whether they still buy cars).
Online retailing AI (e.g., Birchbox)	AI will enable better predictions for what customers want, which may cause firms to move away from a shopping-then-shipping business model and toward a shipping-then-shopping business model.
Fashion-related AI (e.g., Stitch Fix)	AI applications support stylists, who curate a set of clothing items for customers. Stitch Fix's AI analyzes both numeric and image/other non-numeric data.
Sales AI (e.g. Conversica)	AI bots can automate parts of the sales process, augmenting the capabilities of existing sales teams. There may be backlash if customers know (upfront) that they are chatting with an AI bot (even if the AI bot is otherwise capable)
Customer service robots (e.g., Rock'em and Sock'em; Pepper)	Robots with task-automating AI respond to relatively simple customer service requests (e.g., making cocktails).
Emotional support AI (e.g., Replika)	AI aims to provide emotional support to customers by asking meaningful questions, offering social support, and adjusting to users' linguistic syntax.
In-car AI (e.g., Affectiva)	In-car AI that analyzes driver data (e.g., facial expression) to evaluate drivers' emotional and cognitive states.
Customer screening AI (e.g. Kanetix)	AI used to identify customers who should be provided incentives to buy insurance (and avoid those who (1) are already likely to buy and (2) those unlikely to buy).
Business process AI (e.g., IBM Interact)	AI used for multiple (simple) applications, such as customized offers (e.g., Bank of Montreal).
Retail store AI (e.g., Café X, LoweBot, 84.51, Bossa Nova)	Robots that can serve as coffee baristas, respond to simple customer service requests in Lowe's stores, and identifying misshelved items in grocery stores.
Security AI (e.g., Knightscope's K5)	Security robots patrol in offices or malls, equipped with superior sensing capabilities (e.g., thermal cameras).
Spiritual support AI (e.g., BlessU-2; Xian'er)	Customizable robot priest/monk offering blessings in different languages to the user.
Companion robot AI (e.g., Harmony from Realbotix)	Customizable robot companion, which promises reduced loneliness to the user.

2018, p. vi), is “manifested by machines that exhibit aspects of human intelligence” (Huang and Rust 2018, p. 155), and involves machines mimicking “intelligent human behavior” (Syam and Sharma 2018, p. 136). It relies on several key technologies, such as machine learning, natural language processing, rule-based expert systems, neural networks, deep learning, physical robots, and robotic process automation (Davenport 2018). By employing these tools, AI provides a means to “interpret external data correctly, learn from such data, and exhibit flexible adaptation” (Kaplan and Haenlein 2019, p. 17). Another way to describe AI depends not on its underlying technology but rather its marketing and business applications, such as automating business processes, gaining insights from data, or engaging customers and employees (Davenport and Ronanki 2018). We build on this latter perspective. A listing of this research is provided in Table 2.

First, to automate business processes, AI algorithms perform well-defined tasks with little or no human intervention, such as transferring data from email or call centers into recordkeeping systems (updating customer files), replacing lost ATM cards, implementing simple market transactions,

or “reading” documents to extract key provisions using natural language processing. Second, AI can gain insights from vast volumes of customer and transaction data, involving not just numeric but also text, voice, image, and facial expression data. Using AI-enabled analytics, firms then can predict what a customer is likely to buy, anticipate credit fraud before it happens, or deploy targeted digital advertising in real time. For example, stylists working at Stitch Fix, a clothing and styling service, use AI to identify which clothing styles will best suit different customers. The underlying AI integrates data provided by customers' expressed preferences, their Pinterest boards, handwritten notes, similar customers' preferences, and general style trends. Finally, AI can engage customers, before and after the sale. The Conversica AI bot works to move customer transactions along the marketing pipeline, and the AI bot used by 1-800-Flowers provides both sales and customer service support. AI bots offer advantages beyond just 24/7 availability. Not only do these AI bots have lower error rates, but also they free up human agents to deal with more complex cases. Further, AI bot deployment can be scaled up or down as needed, when demand ebbs or flows.

As these descriptions suggest, AI offers the potential to increase revenues and reduce costs. Revenues may increase through improved marketing decisions (e.g., pricing, promotions, product recommendations, enhanced customer engagement); costs may decline due to the automation of simple marketing tasks, customer service, and (structured) market transactions. Furthermore, the above discussions indicate that rather than replacing humans, firms generally are using AI to augment their human employees' capabilities, such as when Stitch Fix uses AI to augment its stylists' efforts to make appropriate choices for clients (Gaudin 2016). This point aligns well with sentiments expressed by Ginni Rometty, the CEO of IBM, who proposed that AI would not lead to a world of man “versus” machine but rather a world of man “plus” machines (Carpenter 2015).

A framework for understanding artificial intelligence

Building on insights from marketing (and more generally business), social sciences (e.g., psychology, sociology), and computer science/robotics, we propose a framework to help customers and firms anticipate how AI is likely to evolve. We consider three AI-related dimensions: levels of intelligence, task type, and whether the AI is embedded in a robot.

Level of intelligence

Task automation versus context awareness Davenport and Kirby (2016) contrast task automation with context awareness. The former involves AI applications that are standardized, or rule based, such that they require consistency and the imposition of logic (Huang and Rust 2018). For example, IBM's Deep Blue applied standardized rules and “brute force” algorithms to beat the best human chess player. Such AI is best suited to contexts with clear rules and predictable outcomes, like chess. On the cruise ship *Symphony of the Seas*, two robots, Rock ‘em and Sock ‘em, make cocktails for customers. Elsewhere, the robot Pepper can provide frontline greetings, and IBM's Watson can provide credit scoring and tax preparation assistance. Notwithstanding that these AI applications involve fairly structured contexts, many firms struggle to implement even these AI applications² and rely on specialized businesses like Infinia ML and Noodle, or consulting firms like Accenture or Deloitte, to develop and set up initial AI initiatives.³

In contrast, context awareness continues to be developed, and researchers in computer science are working on moving

AI capabilities forward, from task automation to context awareness (e.g., Ghahramani 2015; Mnih et al. 2015). Context awareness is a form of intelligence that requires machines and algorithms to “learn how to learn” and extend beyond their initial programming by humans. Such AI applications can address complex, idiosyncratic tasks by applying holistic thinking and context-specific responses (Huang and Rust 2018). However, such capabilities remain distant; a 2016 survey of AI researchers indicated there was only a 50% chance of achieving context awareness (or its equivalent) by 2050 (Müller and Bostrom 2016). Building on the above point, Reese (2018, p. 61) cautions that such AI “does not currently exist... nor is there agreement ... if it is possible.” Nevertheless, this capability constitutes the goal of AI developments, as predicted by compelling examples from science fiction, such as Jarvis from the *Iron Man* movies or Karen from *Spider Man—Homecoming*; both AI can understand new and complex contexts and create solutions therein.

The differences between task automation and context awareness map onto concepts of narrow versus general AI (Baum et al. 2011; Kaplan and Haenlein 2019; Reese 2018). As Kaplan and Haenlein (2019) state, both narrow and general AI may equal or outperform human performance, but narrow AI is focused on a specific domain and cannot learn to extend into new domains, whereas general AI can extend into new domains.

It is important to clarify that although in this paper we consider two levels of intelligence (task automation vs. context awareness), ideally levels of intelligence are best conceptualized as a continuum. Some AI applications have moved beyond task automation but still fall well short of context awareness, such as Google's DeepMind AlphaGo (which beat the world's best Go player), the AI poker player Libratus, and Replika.⁴ These applications represent substantial advances, yet state-of-the-art AI still is closer to task automation (Davenport 2018).

Overview of extant research Research into the psychology of automation (Longoni et al. 2019), examines how customers may respond to AI. Notwithstanding the fact that AI may be more accurate and/ or more reliable than humans, customers have reservations about AI, and these reservations tend to increase as AI moves towards context awareness. In turn, these increased reservations negatively impact the propensity to adopt AI, propensity to use AI, etc. A listing of such research is shown in Table 2. Moving forward, we discuss (separately) issues relating to AI adoption and AI usage.

² Reese (2018, p. 61) cautions that this type of AI is in no way “easy AI.”

³ To clarify, businesses like Infinia ML etc. also provide support moving forward, when the firm initiates more advanced AI initiatives.

⁴ Replika (replika.ai) aims to serve as an AI friend, programmed to ask questions about you and your life that are “meaningful,” and to offer emotional support (French 2018; also see Hassler 2018).

AI adoption Customers appear to hold AI to a higher standard than is normatively appropriate (Gray 2017), as exemplified by the case of driverless cars. Customers should adopt AI if its use leads to significantly fewer accidents; instead, customers impose higher standards and seek zero accidents from AI. Understanding the roots of this excessive caution is important. A preliminary hypothesis suggests that customers trust AI less, and so hold AI to a higher standard, because they believe that AI cannot “feel” (Gray 2017).

Task characteristics also influence AI adoption. To the extent a task appears subjective, involving intuition or affect, customers likely are even less comfortable with AI (Castelo 2019). Research confirms that customers are less willing to use AI for tasks involving subjectivity, intuition, and affect, because they perceive AI as lacking the affective capability or empathy needed to perform such tasks (Castelo et al. 2018).

Tasks differ in their consequences; choosing a movie is relatively less consequential, but steering a car may involve more consequences. Using AI for consequential tasks is perceived as involving more risk, in turn reducing adoption intentions. Early work has found support for this hypothesis, more so among more conservative consumers for whom risks are more salient (Castelo et al. 2018; Castelo and Ward 2016).

Finally, customer characteristics may also impact AI adoption. We build from two points: (1) when outcomes are consequential, this increases perceptions of risk (Bettman 1973), and (2) women perceive more risk in general (Gustafsson 1998) and take on less risk (Byrnes et al. 1999). Hence, early work has found that women (vs. men) are less likely to adopt AI, especially when outcomes are consequential (Castelo and Ward 2016). Moving beyond demographics, other factors also impact the extent of AI adoption, e.g., to the extent a task is salient to a customer’s identity, the customer may be less likely to adopt AI (Castelo 2019). To elaborate, if a certain consumption activity is central to a customer’s identity, then the customer likes to take credit for consumption outcomes (Leung et al. 2018). Some customers perceive that using AI for these consumption activities is tantamount to cheating, and this hinders the attribution of credit post-consumption. Therefore, if an activity is central to a customer’s identity, then the customer may be less likely to adopt AI (for this activity).

AI usage Moving past adoption issues, we note some usage considerations, including how AI should communicate with customers. Customers do not associate AI applications with autonomous goals (Kim and Duhachek 2018); for example, customers do not believe Google’s AlphaGo has the self-driven goal to be a national Go champion. Rather, they believe that this AI application is programmed to play the game Go. Consistent with this perception, customers are more likely to focus on “how” (rather than “why”) the AI application performs; implying that when engaging with AI, customers will be in a low level construal mindset. From extant research, we

know that messages are more effective when the perceived characteristics of the message source and the contents of the actual message match, communication from AI should be more effective when it highlights how rather than why in its messaging (regulatory construal fit; Lee et al. 2009; Motyka et al. 2014). In line with the above, Kim and Duhachek (2018) showed that a message from an AI application is more persuasive when the message is about how to use a product, rather than why to use this product. This is because customers doubt whether AI can “understand” the importance of engaging in certain consumption behaviors.

Next we pivot to factors that impact the propensity of customers to engage with AI. Examining the case of medical decision making, Longoni et al. (2019) show that customers’ reservations are due to their concerns about uniqueness neglect (i.e., the AI is perceived as less able to identify and relate with customers’ unique features). Further, building from prior work (Şimşek and Yalınçetin 2010; also see Haslam et al. 2005), Longoni et al. (2019) show that these reservations are more for customers who have higher scores on the ‘personal sense of uniqueness’ scale. In other work on how customers engage with AI, Luo et al. (2019) examined how (potential) customers engage with AI bots. In reality, AI bots can be as effective as trained salespersons, and 4x as effective as inexperienced salespersons. However, if it is disclosed that the customer is conversing with an AI bot, purchase rates drop by 75%. Linked to points made prior in this paper, because customers perceive the AI bot as less empathetic, they are curt when interacting with AI bots, and so purchase less.

Task type

Task type refers to whether the AI application analyzes numbers versus non-numeric data (e.g., text, voice, images, or facial expressions). These different data types all provide inputs for decision making, but analyzing numbers is substantially easier than analyzing other data forms. Practitioners, such as senior managers from Infinia ML, formulate this categorization slightly differently, noting that data that can be organized into tabular formats are significantly easier to analyze than those data that cannot. In our discussions with employees of Stitch Fix, we gained further clarity on this point. Stitch Fix elicits data from customers using both direct questions about their preferences (which can be put in tabular formats) and indirect elicitations from customers’ Pinterest pages and likes. Stitch Fix uses proprietary AI algorithms to analyze the latter, non-numeric data and regards these data as very useful, because it has learned that customers cannot always articulate their preferences on numeric scales.

The distinction in the above paragraph is critical, because much data is non-tabular in form, and so being able to comprehend and analyze such data significantly enhances the impact of AI. Many AI applications have started to analyze text,

voice, image, and face data inputs. These data inputs are initially in non-numeric formats, but are often translated into numerical formats, e.g., pixel brightness values, relating to images. Applications that can process such data inputs include, for example (1) IPSoft, which processes words spoken to customer agents to interpret what customers want (2) Affectiva, which is working on in-car AI that can sense driver emotion and fatigue and switch control to an autonomous AI, and (3) Cloverleaf's shelfPoint, installed on retail store shelves, which examines customers' facial expressions to analyze their emotional responses at the point of purchase. Although currently AI's abilities to comprehend and analyze such non-numeric data formats remain somewhat limited, developing this ability will be critical for the full realization of the power of AI, and computer scientists are working towards improving AI capabilities in this regard (e.g., LeCun et al. 2015; You et al. 2016).

Separate to the above, it is worth pointing out that the ability to analyze unstructured data may be limited by legacy infrastructures. A senior manager in Infinia ML indicated that often data is stored in formats and structures less amenable to AI deployment. Also, Kroger has an AI application that automates visual inspection of out-of-stock items on its grocery shelves. In an interview with one of the authors of this paper, a Kroger data scientist reported that the proper functioning of Kroger's AI application requires hardware upgrades; specifically, it needs to upgrade its cameras to higher resolution levels if the AI application is to work properly.

AI in robots

Virtuality-reality continuum Most AI is virtual in form. For example, Replika is available on smartphones, and Libratus uses a digital platform. However, AI can also be embedded in a real entity or robot form, with some elements of physical embodiment. The extent to which a form is virtual versus embodied reflects its position on the Milgram virtuality–reality continuum (Milgram et al. 1995). In this sense, researchers and practitioners should conceive of virtual and real forms not as distinct categories but rather as endpoints on a continuum, within which AI entities are spread out. An AI like Conversica is purely virtual, with no physical embodiment—although some companies that use virtual AI do give it names. In contrast, an AI application embedded in a robot barista (e.g., Tippy Robot in Las Vegas) appears somewhere on the continuum between virtuality and reality, because it has some physical embodiment; however, that embodiment can only operate in a narrow range and on a specific task (making a drink). Finally, the AI embedded in proposed multifunctional, companion robots (that today remain under development) would entail substantially more reality, featuring both physical embodiment and the capacity to operate in wide range of contexts (specifically, share physical proximity with

individuals without any protective barrier, travel with individuals, etc.).

Overview of extant research Prior research (Table 2) indicates that using robots offer substantial advantages, especially in cases involving customer interactions. As prior work indicates, customers form more personal bonds with robots than with AI that lack any physical embodiment. For example, individuals enjoy interacting with a physically present robot than with either a robot simulation (on a computer) or a robot presented via teleconference (Wainer et al. 2006). Further, customers empathize with robots. When individuals are asked to administer pain—via electric shocks—to a (physically present) robot or a robot simulation, both of which go on to display marks indicating pain after being subjected to an shocked, individuals empathized more with the physically present robot (Kwak et al. 2013). Finally, customers interacted longer with a robot diet coach than with either a virtually present diet coach or a diet diary in a paper form (Kidd and Breazeal 2008). Other studies find that customers demonstrate reciprocity-based perceptions, e.g., they express more positive perceptions of a care robot that asks for help and then returns this help by offering a favor (Lammer et al. 2014). In a prisoner's dilemma experiment, participants exhibited similar reciprocity levels toward both robot partners and human partners (Sandoval et al. 2016), and their reciprocity towards the robot partner increased even more if the robot provided early signs of cooperation (vs. random behavior). Noting the benefits of embedding AI in robots, work in robotics is examining how best to improve not only the physical capability of robots but also the robot–AI interface (e.g., Adami 2015; Kober et al. 2013; Steels and Brooks 2018). Further, to take advantage of the preference for physical embodiment, some vendors of virtual agents (or bots) try to present these agents as having a physical form. IPsoft's virtual agent, for example, is called Amelia and is often represented by a lifelike avatar image and voice.

However, other research shows that customers' discomfort with AI is accentuated when the AI application is embedded in a robot. As robots appear more humanlike, they become more unnerving, in line with the uncanny valley hypothesis (UVH; Mori 1970).⁵ UVH arises because the appearance of robots “prompts attributions of mind. In particular, we suggest that machines become unnerving when people ascribe to them experience (the capacity to feel and sense), rather than agency (the capacity to act and do)” (Gray and Wegner 2012, p. 125). Such factors may hinder AI adoption.

⁵ Masahiro Mori wrote an influential paper arguing that making robots look more human is beneficial, but only up to a certain point, after which such robots elicit negative reactions. Thus, reactions become negative as robots move. From somewhat human to human-like. Thereafter, if robots look perfectly human, reactions turn positive. The valley reflects these trends, as reactions initially becoming more negative, then turn positive.

Table 2 Select extant research (in the order in which they appear in the paper)

Paper	Domain	Dimension	Takeaways
Agrawal et al. (2018)	BUS		Artificial intelligence (AI) reduces the cost of prediction.
Gans et al. (2017)	BUS		
Rahwan et al. (2019)	CS/R		To best understand AI, bring in insights from not only computer science, but also other disciplines
Shankar (2018)	MKTG		AI “refers to programs, algorithms, systems and machines that demonstrate intelligence” (Shankar 2018, p. vi), is “manifested by machines that exhibit aspects of human intelligence” (Huang and Rust 2018, p. 155), involves machines mimicking “intelligent human behavior” (Syam and Sharma 2018, p. 136), and provides means to “interpret external data correctly, learn from such data, and exhibit flexible adaptation” (Kaplan and Haenlein 2019, p. 17).
Huang and Rust (2018)	MKTG		
Syam and Sharma (2018)	MKTG		Huang and Rust (2018) - Mechanical and analytical intelligences involve simple, rule-based tasks. Intuitive and empathetic intelligences involve complex tasks requiring empathy, holistic thinking and context-specific responses.
Kaplan and Haenlein (2019)	MKTG		Kaplan and Haenlein (2019) – Used the terms narrow versus general AI. Narrow AI somewhat maps onto mechanical and analytical intelligences, whereas general AI maps onto intuitive and empathetic intelligences.
Davenport and Ronanki (2018)	BUS	LVLINT	Another way to describe AI is by stating its marketing and business outcomes, such as automating business processes, gaining insights from data, or engaging customers and employees
Davenport and Kirby (2016)	BUS	LVLINT	Contrasts task automation with context awareness. The former involves AI applications that are standardized, or rule based (akin to narrow AI). The latter is a form of intelligence that requires machines and algorithms to ‘learn how to learn’ and extend beyond their initial programming (akin to general AI).
Ghahramani (2015)	CS/R	LVLINT	How machines can learn from experience, using probabilistic machine learning.
Mnih et al. (2015)	CS/R	LVLINT	How artificial agents can learn to generalize from past experience to new situations, using reinforcement learning.
Müller and Bostrom (2016)	BUS	LVLINT	Artificial general intelligence (AGI) is a hypothetical technology that would be the equivalent of a human intelligence in terms of its flexibility and capability of performing and learning a vast range of tasks (similar to context awareness). In a survey of AI researchers, the median estimate was for a 50% chance of achieving an AGI by 2050 and a 90% chance of achieving one by 2075.
Reese (2018)	BUS	LVLINT	Defines narrow versus general AI and analytical AI versus humanized AI; both contrasts are very similar to the contrast between task automation versus context awareness. Reese (2018) cautions that AGI does not exist, and that there is no guarantee that it ever will.
Baum et al. (2011)	SOC	LVLINT	
Davenport (2018)	BUS	LVLINT	The state-of-the-art AI is closer to task automation than context awareness.
Gray (2017)	PSY	LVLINT	Customers appear to hold AI to a higher standard than is normatively appropriate. A preliminary hypothesis suggests that customers trust AI less, and so hold AI to a higher standard, because they believe that AI cannot “feel”.
Castelo (2019)	MKTG	LVLINT	To the extent a task appears subjective, involving intuition or affect, customers likely are less comfortable with AI (Castelo 2019). Customers are less willing to use AI for tasks involving subjectivity, intuition, and affect, because they perceive AI as lacking the affective capability or empathy needed to perform such tasks (Castelo et al. 2018).
Builds from: Castelo et al. (2018)	MKTG		
Castelo and Ward (2016)	MKTG	LVLINT	Using AI for consequential tasks is perceived as involving more risk, in turn reducing adoption intentions. This is more so amongst (1) conservative consumers, for whom risks are more salient, (2) women, who perceive more risk in general, and take on less risk.
Builds from: Bettman (1973)	MKTG		
Gustafson (1998)	PSY		
Byrnes et al. (1999)	PSY		

Table 2 (continued)

Paper	Domain	Dimension	Takeaways
Leung et al. (2018)	MKTG	LVLINT	If a certain consumption activity is central to a customer's identity, the customer would like to take credit for consumption outcomes. Some customers perceive that using AI for these consumption activities is tantamount to cheating, and this hinders the attribution of credit post-consumption. Hence if an activity is central to a customer's identity, then the customer may be less likely to adopt AI for this activity.
Kim and Duhachek (2018) Builds from: Lee et al. (2009) Motyka et al. (2014)	MKTG MKTG MKTG	LVLINT	Customers do not associate AI applications with autonomous goals (Kim and Duhachek 2018). In line with this perception, customers are more likely to focus on "how" (rather than "why") the AI application performs; implying that when engaging with AI, customers will be in a low level construal mindset. Because persuasion is more effective when the perceived characteristics of the persuasion source and the persuasion message match, communication from AI should be more effective when it highlights how rather than why in its messaging (regulatory construal fit; Lee et al. 2009; Motyka et al. 2014). AI persuasion messages are more effective in persuading consumers to buy the recommended product or services when the message highlights "how" to use the product rather than "why" to use the product. These effects are because customers doubt whether AI can understand "why" it is important for customers to engage in certain behaviors.
Longoni et al. (2019) Builds from: Şimşek and Yalınçetin (2010) Haslam et al. (2005)	MKTG PSY PSY	LVLINT	Examining the case of medical decision making, Longoni et al. (2019) propose that customers' reservations are due to their concerns about uniqueness neglect (i.e., the AI is perceived as less able to identify and relate with customers' unique features). Further, building from prior work (Şimşek and Yalınçetin 2010; also see Haslam et al. 2005), Longoni et al. (2019) propose that these reservations would be more for customers who have higher scores on the 'personal sense of uniqueness' scale (Şimşek and Yalınçetin 2010).
Luo et al. (2019)	MKTG	LVLINT	Examines how (potential) customers engage with AI bots. In reality, AI bots can be as effective as trained salespersons, and 4 times effective as inexperienced salespersons. However, if it is disclosed that the customer is conversing with an AI bot, purchase rates reduce by 75%. Because customers perceive the AI bot as less empathetic, they are curt when interacting with AI bots, and so purchase less. Ties into themes from Castelo et al. (2018).
LeCun et al. (2015)	CS/R	TSKTYPE	How deep learning has improved the state-of-the-art in speech and visual object recognition.
You et al. (2016)	CS/R	TSKTYPE	How using a new algorithm improves visual object recognition.
Milgram et al. (1995)	PSY	ROBOT	Proposes the virtuality-reality continuum.
Wainer et al. (2006)	CS/R	ROBOT	Interacting with a physical robot is perceived as more enjoyable than either interacting with a simulated robot on a computer or interacting with a real robot presented through teleconferencing.
Kwak et al. (2013)	CS/R	ROBOT	When asked to administer electric shocks to a (physical) robot or a simulated robot on a computer screen, individuals empathized more with the (physical) robot.
Kidd and Breazeal (2008)	CS/R	ROBOT	Interactions were longer with a robot diet coach than either a virtual diet coach or a pen-and-paper diet diary.
Lammer et al. (2014)	CS/R	ROBOT	Individuals express reciprocity towards robots.
Adami (2015)	CS/R	ROBOT	With suitable machine learning algorithms, robots can learn from past experiences.
Kober et al. (2013)	CS/R	ROBOT	Reinforcement learning can work for robots embedded with suitable machine learning algorithms.
Mori (1970)	PSY	ROBOT	Making robots look more human is beneficial, but only up to a certain point, after which such robots elicit negative reactions (UVH).
Gray and Wegner (2012)	PSY	ROBOT	Machines are perceived as more unnerving when individuals ascribe to machines the capacity to feel, rather than capacity to do.

Table 2 (continued)

Paper	Domain	Dimension	Takeaways
Mende et al. (2019)	MKTG	ROBOT	Interactions with robots trigger discomfort (linked to UVH) and so further trigger compensatory behaviors.
Boyd and Holton (2018)	SOC	ROBOT	Will the combination of robotics and AI lead to an unprecedented social transformation?
Pedersen et al. (2018)	SOC	ROBOTS	Outlines the issues surrounding use of social robots in medical treatment, care facilities, and private homes. Also, outlines ethical concerns.
André et al. (2018)	MKTG	LVLINT	Because AI facilitates data-driven, micro-targeting marketing offerings, customers should view such offerings favorably, because it reduces search costs. Yet this could undermine customers' perceived autonomy, with implications for their subsequent evaluations and choices.
Aguirre et al. (2015)	MKTG	LVLINT	Proposes the privacy–personalization paradox, whereby individuals balance privacy concerns against the benefits of personalized recommendations.
Wang and Kosinski (2018)	PSY	TSKTYPE	How to use deep neural networks to identify sexual orientation, merely by analyzing facial images

MKTG Marketing; *BUS* Business; *PSY* Psychology; *SOC* Sociology; *CS/R* Computer Science/ Robotics.

Dimension: *LVLINT* levels of intelligence; *TSKTYPE* task type; *ROBOT* whether AI in robots

Moving beyond AI adoption, we pivot to how customers interact with robots with embedded AI. Early research suggests that interactions with AI-embedded robots trigger discomfort (linked to the UVH) and so further trigger (negative) compensatory behaviors, like buying of status goods, or eating more food (Mende et al. 2019). From a theory perspective, this work not only shows the downsides of anthropomorphism (especially in the case of robots), but also the existence of compensatory consumption⁶ specifically linked to robots.

More broadly, sociologists ponder how AI (and specifically robots with embedded AI) might transform economy and society (Boyd and Holton 2018). For example, cloud-based technology facilitates deep learning in robots, which can learn from human agents through repeated interactions. Sociologists particularly note ways that robots may enter multiple aspects of social life, not only in (expected) areas such as service and transportation, but also in domains like the arts and music.

The current state and likely evolution of AI

Short- and medium-term time horizon

In Fig. 1, we combine all the above considerations to depict the current state of AI and its likely evolution. The upper half of Fig. 1 (four cells) relates to task automation and thus the likely state of AI in the short to medium time horizon. The lower half of Fig. 1 (two cells) relates to context awareness applications that are only likely in the long term (if at all), due to the constraints associated with the current state of AI. Note

⁶ Compensatory consumption is consumption “motivated by a desire to offset or reduce a self-discrepancy” (Mandel et al. 2017, p. 134).

that in the lower half of Fig. 1, we do not distinguish between numeric and non-numeric data, because context awareness-capable AI likely will be able to handle any types of data.

The first four use cases, associated with short to medium term developments, involve task automation (see Fig. 1).

Cell 1: Controller of numerical data The first cell in Fig. 1 reflects what AI can do very well, namely, statistical analyses of numeric data using machine learning. A typical use case is the application of AI to optimize prices (Antonio 2018). Pricing strategies must balance two competing concerns; that the price is low enough to attract customers versus high enough to enable the firm to earn sufficient profits. Firms use AI to analyze vast amounts of numeric data (including less intuitive predictor variables) to both set optimal prices and then change prices in real time. For example, Kanetix helps Canadian customers find deals on car insurance by allowing prospective buyers to compare and evaluate policies and rates offered by more than 50 providers. Scott Emberley, the Business Development Director of integrate.ai, which partnered with Kanetix to build an AI application, indicated that the goal was identify three sets of customers (1) those highly likely to buy, (2) those very unlikely to buy, and (3) those in-between. Thereafter, Kanetix would direct their advertising towards these “in-between” customers, which would provide the greatest returns, and not expend efforts on those either very likely to buy or very unlikely to buy. With four years of data, integrate.ai developed a machine learning model that could identify such customers. Five months later, Kanetix estimated 2.3 times return on its AI investment, and a more than 20% increase in sales among previously undecided customers. In another example, the Bank of Montreal (BMO) uses IBM Interact to analyze customer data across all its channels and identify personalized product offerings. If a customer

has been exploring mortgages on BMO’s site and later calls the contact center, IBM Interact prioritizes the list of available mortgage offers for the contact center service agent—in effect augmenting agents’ capabilities and facilitating more relevant customer conversations.

Cell 2: Controller of data Efforts to analyze non-numeric data offer the potential to improve understanding of what customers want, and firms’ customer service. Some AI applications can analyze non-numeric data (in some cases, after conversion to numeric data), primarily using speech and image recognition capabilities achieved with deep learning neural networks (Chui et al. 2018). For example, Conversica AI, as manifested in a virtual AI assistant named Angie, sends out-bound emails to up to 30,000 leads per month, then interprets the responses to identify the most promising leads (Power 2017). Angie engages in initial conversation with the prospect, and then routes to most promising leads to a (human) salesperson. In effect, Conversica’s AI augments salespersons’ capabilities. In a pilot test with a telecommunications company called Century Link, Angie appropriately understood more than 95% of emails received (and sent the rest to human agents for interpretation), and Century Link earned a 20-fold return on its investments in Angie.

The Stitch Fix’s business model offers another example. As we noted, Stitch Fix delivers apparel directly to customers (Wilson et al. 2016), without requiring the customers to actually engage in a formal shopping task. No Stitch Fix retail

location exists. Instead, customers fill out style surveys, provide their physical measurements, evaluate sample styles, create links to their Pinterest boards, and send in personal notes. As may be expected, customers have trouble explicating their exact style preferences using words and numbers, but their pins and likes can be (better) indicators of their preferences. Stitch Fix’s proprietary machine learning algorithms examine numbers, words, and Pinterest pins, then summarize the findings for the company’s fashion stylists, who in turn select suitable clothing to send to each customer. The above example illustrates the need to suitably balance AI input and human input; senior managers from Stitch Fix told us that—in their experience—their AI works best when it augments the (human) stylists’ capabilities.

Noting that the AI applications in companies like Conversica and Stitch Fix use all types of data (i.e., use numeric data and non-numeric data), we term the AI applications in this cell as reflecting “Controller of Data.”

Cell 3: Numerical data robot This cell is similar to cell 1, except that it incorporates AI embedded in a robotic form, and so these AI applications can best be described as robots that process numerical data inputs. Such robots are well suited to retail environments with well-structured operations. At Café X, for example, a robot barista can serve up to 120 coffees per hour (Hochman 2018). Each robotic barista features a \$25,000, six-axis animatronic arm. Customers place orders on a kiosk touchscreen (or via an app), so all inputs

Fig. 1 AI framework

Task automation technologies, deployed currently or to be deployed in the short to medium term

	Digital form	Robot form
Analyze numbers	1 – Controller of Numerical Data Business Use Case Kanetix IBM	3 – Numerical Data Robot Business Use Case Café X Tipsy Robot
Analyze text, voice, faces, images	2 – Controller of Data Business Use Case Conversica Stitch Fix Replika	4 – Data Robot Business Use Case LoweBot 84.51/ Kroger Walmart/ Bossa Nova K5 from Knightscape

Context awareness technologies that *may* be deployed in the long term

	Digital form	Robot form
Analyze numbers, text, voice, faces, image	5 – Data Virtuoso Example Use Case Jarvis	6 – Robot Expert Example Use Case Dorian

are numeric. As in a regular coffeehouse, customers can select various options: latte or espresso, with different amounts of froth, and various ingredients such as organic Swedish oat milk. The goal is not to replace baristas, but rather to augment baristas' capabilities by taking over more routine operations. The Cafe X robot barista augments the capabilities of the human barista, who can then focus on providing high-quality customer service, and also facilitating what the company calls "coffee education" (e.g., managing tastings).

Cell 4: Data robot This cell is similar to cell 2, except that the robotic form can process all types of data (not just numeric data). For example, the LoweBot at Lowe's Home Improvement stores (Hullinger 2016) can scan a product held up by a customer (or listen to the customer speak the name of the desired product), confirm whether the item is in stock, and then roll along with the customer to the exact spot in the store where he or she can find the product. This task requires comprehension and examination of both numeric and non-numeric data, as well as an indoor navigation capability, which represents a significant advance over the capabilities embodied in the Café X robot. Using the LoweBot augments the capabilities of Lowe's human sales associates, allowing focus on more complex customer service requests.

Other retailers have similar applications. Our discussions with senior managers at 84.51⁷ indicate that they are working with Kroger to implement in-store robots that can identify misshelved or out-of-stock items. In another example, Walmart has partnered with Bossa Nova Robotics to deploy robots in its stores to scan shelves. The goal appears to be to get robots to perform tasks that repeat and are predictable, enabling (human) associates to focus on serving customers (Avalos 2018).

Finally, security robots, such as the K5 from Knightscope, roam offices and malls at night. These robots have better sensing capabilities than humans, because they incorporate thermal cameras and other high-technology sensing tools. Here again, the objective is to augment human security guards' capabilities (Robinson 2017).

Long-term time horizon

For completeness, we also examine what might happen when AI applications incorporate context awareness, as summarized in the two cells in the lower half of Fig. 1. We reiterate that there is no indication that such developments will occur in the short or medium term, as exemplified by the case of driverless cars. Tesla has removed any "self-driving" labels from its website, noting that these labels were causing confusion (Hawkins 2019). The CEO of Waymo admits that driverless cars are unable to drive in poor weather conditions without

human input (Lashinsky 2019). Put simply, the dream of getting into a driverless car outside in one city, falling asleep, and waking up in another city is not reality and may not be achieved anytime soon. Even the less consequential forms of AI remain problematic. Google's AlphaGo Zero might have successfully learned the complex game of Go in a short period, using adversarial networks that pit two (competing) AI systems against each other so that they can learn; yet in this case, the outcome space was very well defined. Furthermore, all these AI systems received significant training data. In contrast, the outcome spaces (i.e., business domains) for most likely AI applications are poorly defined, and relevant training data is hard to obtain. These points reiterate the challenges of moving from task automation to context awareness. As such, the use cases we present for the last two cells are hypothetical, and this section is deliberately brief, reflecting that our discussion is more aspirational than descriptive of any near-term reality.

Cell 5: Data virtuoso Advanced AI could be embedded in a digital form, as exemplified by the AI Jarvis in *Iron Man* movies. Jarvis has advanced data capabilities that can examine multiple data types. Perhaps most notably, Jarvis adapts to new contexts, beyond those for which it has been trained, such as when it hides from the more advanced AI Ultron and finds ways to thwart Ultron's hacking attempts. Futurists would have us believe that such AI will emerge in the long term, with strong predictive abilities for customers' preferences and high capability levels for managing customer service. Thus, the term virtuoso seems appropriate for such AI.

Cell 6: Robot experts An advanced AI also could be embedded in a robot form, such as the AI Dorian from the television show *Almost Human*. Dorian's advanced capabilities include facial recognition, bio scans, analyses of non-numeric stimuli such as DNA, speed-reading, speaking multiple languages, and taking the temperature of fluids using his finger. Like Jarvis, Dorian can adapt to a variety of new contexts. Futurists predict that such robot experts will emerge in the long term to serve as companions that meet various customer needs (e.g., in-home service, home security, medical support). Such robots even might be able to bond emotionally with (human) customers, and potentially replace human partners and animal partners.

Agenda for future research

Having described AI and presented a framework to better understand it, we pivot to outlining some important areas for future research. These include how firms may need to change their marketing strategy, how customers' behaviors will be impacted, and issues relevant to policymakers. We outline these areas in Fig. 2, linking these to the cells in Fig. 1.

⁷ This consulting firm is a subsidiary of Kroger and provides retail insights to Kroger and its partners; it has strong analytics and AI capabilities.

AI and marketing strategy

Predictive ability Because AI can help firms predict what customers will buy, using AI should lead to substantial improvements in predictive ability. Contingent on levels of predictive accuracy, firms may even substantially change their business models, providing goods and services to customers on an ongoing basis based on data and predictions about their needs. Multiple research opportunities thus emerge, related to different customer purchase behaviors and marketing strategies. One especially important research area may relate to how well prediction AI-driven algorithms may extend to forecasting demand for really new products (RNPs; described in Zhao et al. 2012). AI algorithms probably have good predictive ability for incrementally new products; the open question is whether they will have good predictive ability for RNPs. For AI algorithms to do so would presumably require data on RNPs that would be used in training machine learning models; this is often not readily available. Further, when examining

how best to make predictions for RNPs, research can also examine how best to combine AI-driven insights with human judgment.

AI is expected to play an important role in predicting not only what customers want to buy, but also what price to charge, and whether price promotions should be offered (Shankar 2018). Price and price promotions are important drivers of sales (Biswas et al. 2013; Guha et al. 2018), and so are an important area of research for marketing researchers. Thus, an important area for future research relates to how AI can be best used to predict what prices are optimal and whether or not price promotions should be offered.

Another important research avenue pertains to allocations of advertising resources. Much advertising focuses on developing customer awareness and driving customers' information search. Would these advertising dollars be required in the future, wherein firms may be able to better predict customers' preferences, and thus would not need to advertise as much?

Fig. 2 Research agenda for AI. Notes: As noted in the text, the sales AI application will be more effective if it can process both numeric and non-numeric data, and hence is more related to the Controller of Data cell. This is more likely for more advanced robots, and so more likely to be relevant to robots able to handle non-numeric data (notably voice), and hence more related to perhaps the Data Robot cell, but more so to the Robot Expert cell

Task automation

Context awareness

	Digital form	Robot form
Analyze numbers	1 – Controller of Numerical Data Research Agenda <i>Predictive ability (MS)</i> <i>AI adoption (CB)</i> – <i>negative response to AI</i> – <i>state and trait moderators</i> <i>AI usage (CB)</i> – <i>primed mindset</i> – <i>Post AI issues (CB)</i> – <i>perceived loss of autonomy</i> – <i>state and trait moderators</i> <i>Data privacy (P)</i> <i>Bias (P)</i> <i>Ethics (P)</i>	3 – Numerical Data Robot Research Agenda <i>similar to Controller of Numerical Data cell</i> + <i>Affective responses to robots (CB)</i>
Analyze text, voice, faces, images	2 – Controller of Data Research Agenda <i>similar to Controller of Numerical Data cell</i> + <i>Sales (MS)</i> <i>AI adoption for spiritual well-being (CB)</i>	4 – Data Robot Research Agenda <i>similar to Controller of Data cell</i> + <i>UVH (CB)</i> <i>Loss of human connectedness (CB)</i>
Analyze numbers, text, voice, faces, image	5 – Data Virtuoso Research Agenda <i>similar to Controller of Data cell</i>	6 – Robot Expert Research Agenda <i>similar to Data Robot cell</i>

Notes: MS = marketing strategy; CB = consumer behavior; P = public policy.

Sales and AI As we discussed with regard to Conversica, AI may alter all stages of the sales process, from prospecting to pre-approach to presentation to follow-up (Singh et al. 2019; Syam and Sharma 2018). Thus, a wide variety of research questions arise:

- Can AI analyze customer communication and other customer information (e.g., social media posts) in ways to devise future communications that are more persuasive or increase engagement?
- Can AI provide real-time feedback to salespeople to help them improve their sales pitches, based on assessments of customers' verbal and facial responses?
- How might AI combine text and other communication inputs (e.g., voice data), actual customer behavior, and other information (e.g., behaviors of similar customers) to predict repurchases? This effort demands non-numeric data, in line with cells 2, 4, 5 and 6.
- Considering Luo et al.'s (2019) findings, how should firms deploy AI sales bots effectively?

Answering these questions could help firms design sales to take the most advantage of AI.

In addition, firms need to consider how they (re)organize their sales and innovation processes. These points are not listed in Fig. 2, as they do not tie neatly into the cells shown in Fig. 1.

Sales process In the presence of AI, how should sales be organized and what skills will salespeople need? First, how best to structure the sales organization wherein organizational components include both AI bots and human salespeople. Secondly, how should the firm manage the tradeoff between AI focusing on customers' expressed needs versus salespeople being relatively better able to manage issues like customer stewardship. Lastly, will salespeople be able to be trained/ to be able to manage customers' concerns relating to AI, specifically issues related to data privacy and ethics. It is clear that sales processes will require innovation related not only to AI technologies, but also in job design and skills (Barro and Davenport 2019).

AI innovation process Because the impact of AI is uncertain, firms need to figure out how best to (continually) develop AI. In our discussions with senior managers at Stitch Fix, they indicated that the company encourages its data scientists to pursue projects on their own (Colson 2018), such that they continually engage in preliminary testing of new project ideas. One Stitch Fix data scientist created a Tinder-like app called Style Shuffle, to allow users to indicate preferences for various clothing styles. This app not only informed stylists about customers' preferences (the expected benefit) but also helped match stylists with specific customers (an unexpected benefit).

Clothing suggestions from stylists who “swiped” on the app similarly to particular customers elicited more positive responses from the customers (i.e., both qualitative feedback about the stylist and increased sales of clothes curated by that stylist). When implementing AI, firms thus may achieve better outcomes if they let their data scientists spend some amount of time on unauthorized “pet projects,” a research and development practice already in place in firms like 3 M (Shum and Lin 2007). Researching the best way to implement AI, to take advantage of both expected and unexpected benefits, is a fruitful area for research.

Modeling the evolution of AI Finally, firms need to develop realistic expectations, because “in the short run, AI will provide evolutionary benefits; in the long run, it is likely to be revolutionary” (Davenport 2018, p. 7). That is, the benefits of AI could be overestimated in the short term but underestimated in the long term, a point (sometimes called Amara's Law) in accordance with Gartner's hype cycle model of how new technologies evolve (Dedehayir and Steinert 2016; also see van Lente et al. 2013; Shankar 2018). This view is popular among practitioners, according to our personal discussions and interviews with various senior managers. Will the evolution of AI reflect this model, or will its evolution differ and more closely map onto models that also integrate more traditional innovation models (e.g., Roger's model, the Bass model)? Research that tests which innovation model best predicts AI evolutions will be useful.

AI and customer behavior

New technologies often alter customer behavior (e.g., Giebelhausen et al. 2014; Groom et al. 2011; Hoffman and Novak 2018; Moon 2003), and we expect that AI will do so as well. We propose three research topics, related to AI adoption, AI usage, and post-adoption issues.

AI adoption As a general point, due to a wide variety of factors, customers view AI negatively, which is a barrier to adoption. As noted, these negative views often stem from customers' sense that AI is unable to feel (Castelo et al. 2018; Gray 2017) or that AI is relatively less able to identify what is unique about each customer (Longoni et al. 2019). Also Luo et al. (2019) suggest that customers perceive AI bots as being less empathetic. Customers also are less likely to adopt AI in consequential domains (Castelo et al. 2018; Castelo and Ward 2016) and for tasks salient to their identity (Castelo 2019; Leung et al. 2018).

Thus, an important area for future research, important from the standpoint of both research and practice, would be to examine how best to mitigate the impact of the above. Initial brainstorming with fellow researchers and with practitioners suggests that positioning AI as a learning (artificial) organism,

or else positioning the AI application as one that combines AI and human inputs (as in Stitch Fix), may help partially mitigate the impact of the points above. Longoni et al. (2019) propose that offering customers the opportunity to slightly modify the AI may get these customers to look past uniqueness neglect, and focus more on the benefits of personalization. This too may be a way to mitigate the points raised prior.

The discomfort with AI is accentuated in case the AI application is embedded in a robot. As robots become more humanlike, then due to the UVH, customers find these robots unnerving. Such factors may hinder AI adoption and deserve study. An interesting moderator of this effect—worth investigating—may be whether the AI form is perceived by customers as a servant or partner; UVH effects may be stronger if AI achieves partner status. Also deserving of study is other ways of mitigating such effects. Early efforts in this direction involve trying to prime empathy, by convincing customers that robots have some ability to see things from the customers' viewpoint, and (also) have some ability to feel sympathy for the customer if the customer were suffering (Castelo 2019). Other possible methods could relate to anthropomorphizing the AI, as this may persuade customers that the AI has somewhat more empathy (this point needs to balance with concerns about the UVH).

Sociologists appear especially interested in how robots with embedded AI might make inroads into society (Boyd and Holton 2018, p. 338), noting that “complexities arise when cultural preferences associated with human as opposed to machine delivery of personal services are considered. Do ... consumers find social robots acceptable?” Broadly speaking, research can address how attitudes toward robots vary by culture (Li et al. 2010). Beyond concerns associated with culture, it may be pertinent to examine which other trait factors determine whether customers are willing to have their hair styled by robots or accept childcare/elderly care services delivered by robots (Pedersen et al. 2018). In addition to physical well-being considerations, some sociologists suggest that robots may assist with spiritual well-being (Fleming 2019), as exemplified by the robot priest BlessU-2 (Sherwood 2017) and Buddhist monk Xian'er (Andrews 2016). Understanding how robots with embedded AI can assist in various ways, beyond improving customers' physical well-being, is a good area for research.

AI usage When customers interact with an AI application, it might prime a low-level construal mindset (Kim and Duhachek 2018). Research should determine what other mindsets might be primed by AI, e.g., AI may prime prevention focus among customers for whom AI is a relatively new technology. Related insights would have implications for how the AI application should communicate with the customer, because communication exerts stronger impacts when it fits with (a primed) mindset.

When AI is embedded in robots, the robots likely have important roles in customers' lives, functioning as frontline service providers (Wirtz et al. 2018), companions, nannies, or pet replacements. In addition to the UVH-related challenges documented previously, some research results suggest that interactions with AI-embedded robots trigger discomfort and compensatory behaviors (Mende et al. 2019). It is important to determine when customers perceive AI-embedded robots negatively and whether these perceptions may improve over time.

Finally, if customers' ideal preferences actually differ from their past behaviors (e.g., customers trying to stop eating unhealthy foods), AI might make it harder for them to find and move toward their preferred options, by only presenting them with choices reflecting their past behaviors. The widespread “retargeting” of digital ads is one example of this phenomenon. How to train AI to best manage this issue?

Post-adoption The downstream consequences of AI adoption also suggest some relevant research topics. In particular, customers might perceive a loss of autonomy if AI can substantially predict their preferences. In theory, because AI facilitates data-driven, micro-targeting marketing offerings (e.g., Gans et al. 2017; Luo et al. 2019), customers should view offerings more favorably, because it reduces their search costs. Yet it also could undermine customers' perceived autonomy, with implications for their evaluations and choices (André et al. 2018). If customers learn that an AI algorithm can predict their preferred choices, they may deliberately choose a non-preferred option, to reaffirm their autonomy (André et al. 2018; Schrift et al. 2017). Such considerations evoke a variety of research questions. For example, which factors determine whether (and how much) customers value perceived autonomy in AI-mediated choice settings? In this regard, it may be helpful to examine individual difference variables, such as culture and whether customers regard AI as a servant or partner. Research also might address state factors, such as the product type; perceived autonomy may be less relevant for utilitarian product choices than for hedonic ones, because of differential links to customers' identity.

Also, there is a generalized fear of a loss of human connectedness, if humans form bonds with robots with embedded AI. The popular press (e.g., Marr 2019) stokes concerns about robots with embedded AI becoming popular (over humans) as partners. Robots like Harmony (by Realbotix) appear promising in this regard, able to take on different personae and display some expression. But such robots arguably could be damaging to society at large, by increasing social isolation, reducing the incidence of marriage, or reducing birthrates—which is critical for countries like Japan, where birthrates are already low. This point suggests some interesting research opportunities.

AI and policy issues

Finally, AI is of interest to policymakers. We note three broad areas in which policymakers seek to ensure that firms strike a suitable balance between their own commercial interests and the interests of customers: data privacy, bias, and ethics.⁸ All three of these areas generally align with all cells in Fig. 2.

Data privacy Today, the combination of AI and big data implies that firms know much about their customers (Wilson 2018). Hence, two issues deserve research attention. First, customers worry about the privacy of their data (Martin and Murphy 2017; Martin et al. 2017). Privacy is complicated (Tucker 2018), for three reasons: (1) the low cost of storage implies that data may exist substantially longer than was intended, (2) data may be repackaged and reused for rationales different than those intended, and (3) data for a certain individual may contain information about other individuals. Policy related to data privacy requires balancing two competing priorities. Too little protection means that customers may not adopt AI-related applications; too much regulation may strangle innovation.

Second, important research questions pertain to whether data privacy management efforts should be driven by legal regulations or self-regulation, in that “it is not clear yet if market driven incentives will be sufficient for firms to adopt policies that favor consumers or whether regulatory oversight is required to ensure a fair outcome for consumers” (Verhoef et al. 2017, p. 7). Cultural perspectives on data privacy also vary, which is an important consideration; some have suggested that the lack of data privacy in China, for example, is consistent with Confucian cultural ideals (Smith 2019).

Third, we need insights into how best to acknowledge and address privacy concerns at the moment data are collected, as well as how to manage data privacy failures (e.g., data breaches). Amazon already sells doorbells with cameras (the Ring device) and may be planning to add facial identification AI to the devices (Fowler 2019). Customers thus may become concerned if Amazon has access to data recorded through Ring, which it could use or sell. Neighbors also might protest if Ring cameras record their front yard activities without their permission. Also, the data from Ring arguably could be subpoenaed by law enforcement agencies or obtained illegally by hackers. Such issues reflect topics for further research.

Finally, we consider the privacy–personalization paradox (Aguirre et al. 2015). Customers must balance privacy concerns against the benefits of personalized recommendations and offers. Important questions relate to how customers determine the optimal trade-off, including which individual difference variables and state variables might moderate their

choices. Does the trade-off depend on the product category or the level of the customer’s trust in the firm, for example? Also, how would this trade-off shift over time?

Bias The potential algorithmic bias embedded in AI applications could stem from multiple causes (Villasenor 2019), including the data sets that inform AI. For example, Amazon abandoned a tool that used AI intelligence to rate job applicants, in part because it discriminated against female applicants (Weissman 2018). This bias emerged because the training data sets used to develop the algorithm were based on data relating to previous applicants, who were predominantly men. Exacerbating the issue, many AI algorithms are opaque black boxes, so it is difficult to isolate which exact factors these algorithms consider. Testing for whether there is bias in AI applications is an important topic.

In addition, AI may not be able to distinguish attributes that could induce potential bias. Villasenor (2019) argues as follows. In general, it may not be offensive when insurance companies treat men and women differently, with one set of premiums for male drivers and another set of premiums for female drivers. Does this mean that it is acceptable for AI to calculate auto insurance premiums based on religion? Many would argue that basing auto insurance premiums on religion is offensive, but from the point of view of an AI algorithm designed to “slice-and-dice” data in every way possible, the distinction between using gender versus religion, as a basis to determining auto insurance rates, may not be obvious. All this suggests that issues relating to bias remain a non-trivial problem (Knight 2017).

Ethics Finally, AI developers must grapple with ethics; we highlight two issues. First, data privacy choices may reflect a firm’s strategy (e.g., if it wants to be perceived as a trusted firm; Martin and Murphy 2017; also see Goldfarb and Tucker 2013) but also could be driven by ethical concerns. In this sense, research should address “how can normative ethical theory pave the way for what organizations should be doing to exceed consumer privacy expectations, as well as to over-comply with legal mandates in order to preserve their ability to self-regulate” (Martin and Murphy 2017, p. 152). A related research topic might involve examining how ethical concerns about AI vary across cultures.

Second, firms choose to deploy AI by defining which problems the AI will tackle. For example, two Stanford researchers used deep neural networks to identify people’s sexual orientation, merely by analyzing facial images (Wang and Kosinski 2018). The deep neural network tools (vs. human judges) were better able to differentiate between gay and straight men. However, the work raised ethical concerns, in that many argued that this AI-based technology may be used by spouses on their partners (if they suspected their partners were closeted), or—more frighteningly—may be used by certain governments to “out” and then prosecute certain populations (Levin

⁸ Firms are aware of this, and are taking steps to suitably respond (Deloitte Insights, as reported in Schatsky et al. 2019)

2017). An important topic for research thus is to address up-front the types of applications for which AI should be used for (or, should not be used for).

Conclusion

This paper outlines a framework to understand how AI will impact the future of marketing, specifically to outline how AI may influence marketing strategies and customer behaviors. We build on prior work, as well as build from extensive interactions with practitioners. First, we develop a multidimensional framework for the evolution of AI, noting the importance of dimensions pertaining to intelligence levels, task types, and whether the AI is embedded in a physical robot. In so doing, we provide the first attempt to integrate all three dimensions in a single framework. We also make two (cautionary) points. First, the short to medium term impacts of AI may be more limited than the popular press would suggest. Second, we suggest that AI will be more effective if it is deployed in ways that augment (rather than replace) human managers.

To examine the full scope of the impact of AI, we propose a research agenda covering three broad areas: (1) how firms' marketing strategies will change, (2) how customers' behaviors will change, and (3) issues related to data privacy, bias, and ethics. This research agenda warrants consideration by academia, firms, and policy experts, with the recognition that although AI already has had some impact on marketing, it will exert substantially more impact in the future, and so there is much still to learn. We hope that this research agenda motivates and guides continued research into AI.

Acknowledgements The authors thank Eric Colson (Stitch Fix), James Kotecki (Infinitia ML) and Charu Kohli (CognitiveScale), who shared their time, materials and insights with the authors as this paper developed. Thomas Davenport and Dhruv Grewal dedicate this article to the memory of their dear friend and colleague, Dean Bala Iyer of Babson University.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Adami, C. (2015). Artificial intelligence: Robots with instincts. *Nature*, 521(7553), 426–427.
- Agrawal, A., Gans, J. S., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business School Press.
- Aguiar, E., Mahr, D., Grewal, D., de Ruyter, K., & Wetzels, M. (2015). Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of Retailing*, 91(1), 34–49.
- André, Q., Carmon, Z., Wertenbroch, K., Crum, A., Frank, D., Goldstein, W., et al. (2018). Consumer choice and autonomy in the age of artificial intelligence and big data. *Customer Needs and Solutions*, 5(1–2), 28–37.
- Andrews, T. (2016). Meet the robot monk spreading the teachings of Buddhism around China. *Washington Post*, April 27. Retrieved February 11, 2019 from https://www.washingtonpost.com/news/morning-mix/wp/2016/04/27/meet-the-robot-monk-spreading-the-teachings-of-buddhism-around-china/?utm_term=.fed52d90bfff3. Accessed 11 Feb 2019
- Antonio, V. (2018). How AI is changing sales. *Harvard Business Review*, July 30. Retrieved February 11, 2019 from <https://hbr.org/2018/07/how-ai-is-changing-sales>.
- Avalos, G. (2018). Walmart tests shelf-scanning robots in Bay Area. *The Mercury News*, March 20. Retrieved February 11, 2019 from <https://www.mercurynews.com/2018/03/20/walmart-tests-shelf-scanning-robots-bay-area/>.
- Barro, S., & Davenport, T. H. (2019). People and machines: Partners in innovation. *MIT Sloan Management Review*, 60(4), 22–28.
- Baum, S. D., Goertzel, B., & Goertzel, T. G. (2011). How long until human-level AI? Results from an expert assessment. *Technological Forecasting and Social Change*, 78(1), 185–195.
- Bettman, J. (1973). Perceived risk and its components: A model and empirical test. *Journal of Marketing*, 10(2), 184–190.
- Biswas, A., Bhowmick, S., Guha, A., & Grewal, D. (2013). Consumer evaluations of sale prices: Role of the subtraction principle. *Journal of Marketing*, 77(4), 49–66.
- Boyd, R., & Holton, R. J. (2018). Technology, innovation, employment and power: Does robotics and artificial intelligence really mean social transformation? *Journal of Sociology*, 54(3), 331–345.
- Burrows, L. (2019). The science of the artificial. Retrieved June 12, 2019 from <https://www.seas.harvard.edu/news/2019/05/science-of-artificial>.
- Byrnes, J. P., Miller, D. C., & Schafer, W. D. (1999). Gender differences in risk taking: A meta-analysis. *Psychological Bulletin*, 125(3), 367.
- Carpenter, J. (2015). IBM's Virginia Rometty tells NU grads: Technology will enhance us. Retrieved February 11, 2019 from <https://www.chicagotribune.com/bluesky/originals/ct-northwestern-virginia-rometty-ibm-bsi-20150619-story.html>.
- Castelo, N. (2019). *Blurring the line between human and machine: Marketing artificial intelligence* (doctoral dissertation). Retrieved from Columbia University Academic Commons. <https://doi.org/10.7916/d8-k7vk-0s40>.
- Castelo, N., & Ward, A. (2016). Political affiliation moderates attitudes towards artificial intelligence. In P. Moreau & S. Puntoni (Eds.), *NA - advances in consumer research* (pp. 723–723). Duluth, MN: Association for Consumer Research.
- Castelo, N., Bos, M., & Lehman, D. (2018). *Consumer adoption of algorithms that blur the line between human and machine*. Graduate School of Business: Columbia University Working Paper.
- Chui, M., Manyika, J., Miremadi, M., Henke, N., Chung, R., Nel, P., & Malhotra, S. (2018). Notes from the AI frontier: Applications and value of deep learning. McKinsey global institute discussion paper, April 2018. Retrieved June 12, 2019 from <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning>.
- Colson, E. (2018). Curiosity-driven data science. *Harvard Business Review*, November 27. Retrieved February 11, 2019 from <https://hbr.org/2018/11/curiosity-driven-data-science>
- Columbus, L. (2019). 10 charts that will change your perspective of AI in marketing. *Forbes*, July 07. Retrieved on July 09, 2019 from <https://www.forbes.com/sites/louiscolombus/2019/07/07/10-charts-that-will-change-your-perspective-of-ai-in-marketing/>

- Davenport, T. H. (2018). *The AI advantage: How to put the artificial intelligence revolution to work*. MIT Press.
- Davenport, T. H., & Kirby, J. (2016). Just how smart are smart machines? *MIT Sloan Management Review*, 57(3), 21–25.
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
- Davenport, T. H., D'Alle Mule, L., & Lucker, J. (2011). Know what your customers want before they do. *Harvard Business Review*, 89(12), 84–92.
- Dedehayir, O., & Steinert, M. (2016). The hype cycle model: A review and future directions. *Technological Forecasting and Social Change*, 108, 28–41.
- Fleming, P. (2019). Robots and organization studies: Why robots might not want to steal your job. *Organization Studies*, 40(1), 23–38.
- Fowler, G. (2019). The doorbells have eyes: The privacy battles brewing over home security cameras. Retrieved February 11, 2019 from <https://www.scribd.com/news/business/2019/02/01/doorbells-have-eyes/>.
- French, K. (2018). Your new best friend: AI chatbot. Retrieved February 13, 2019 from <https://futurism.com/ai-chatbot-meaningful-conversation>.
- Gans, J., Agrawal, A., & Goldfarb, A. (2017). How AI will change strategy: A thought experiment. Harvard business review online. Retrieved February 11, 2019 from <https://hbr.org/product/how-ai-will-change-strategy-a-thought-experiment/H03XDI-PDF-ENG>.
- Gaudin, S. (2016). At stitch fix, data scientists and a.I. become personal stylists. Retrieved February 11, 2019 from <https://www.computerworld.com/article/3067264/artificial-intelligence-at-stitch-fix-data-scientists-and-ai-become-personal-stylists.html>.
- Ghahramani, Z. (2015). Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553), 452–459.
- Giebelhausen, M., Robinson, S. G., Sirianni, N. J., & Brady, M. K. (2014). Touch versus tech: When technology functions as a barrier or a benefit to service encounters. *Journal of Marketing*, 78(4), 113–124.
- Goldfarb, A., & Tucker, C. (2013). Why managing consumer privacy can be an opportunity. *MIT Sloan Management Review*, 54(3), 10–12.
- Gray, K. (2017). AI can be a troublesome teammate. Harvard Business Review, July 20. Retrieved February 11, 2019 from <https://hbr.org/2017/07/ai-can-be-a-troublesome-teammate>.
- Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition*, 125(1), 125–130.
- Groom, V., Srinivasan, V., Bethel, C. L., Murphy, R., Dole, L., & Nass, C. (2011, May). Responses to robot social roles and social role framing. In *Proceedings of the IEEE International Conference on Collaboration Technologies and Systems (CTS)* (pp. 194–203).
- Guha, A., Biswas, A., Grewal, D., Verma, S., Banerjee, S., & Nordfält, J. (2018). Reframing the discount as a comparison against the sale price: Does it make the discount more attractive? *Journal of Marketing Research*, 55(3), 339–351.
- Gustafsson, P. E. (1998). Gender Differences in risk perception: Theoretical and methodological perspectives. *Risk Analysis*, 18(6), 805–811.
- Harding, K. (2017). AI and machine learning for predictive data scoring. Retrieved February 11, 2019 from <https://www.objectiveit.com/blog/use-ai-and-machine-learning-for-predictive-lead-scoring-on-13-February-2019>.
- Haslam, N., Bain, P., Douge, L., Lee, M., & Bastian, B. (2005). More human than you: Attributing humanness to self and others. *Journal of Personality and Social Psychology*, 89(6), 937–950.
- Hassler, C. A. (2018). Meet Replika, the AI Bot that wants to be your best friend. Retrieved February 11, 2019 from <https://www.popsugar.com/news/Replika-Bot-AI-App-Review-Interview-Eugenia-Kuyda-44216396>.
- Hawkins, A. (2019). No, Elon, the navigate on autopilot feature is not ‘full self driving’. Retrieved February 11, 2019 from <https://www.theverge.com/2019/1/30/18204427/tesla-autopilot-elon-musk-full-self-driving-confusion>.
- Hayes, A. (2015). The unintended consequences of self-driving cars. Retrieved February 11, 2019 from <https://www.investopedia.com/articles/investing/090215/unintended-consequences-selfdriving-cars.asp>.
- Hochman, D. (2018). This \$25,000 robotic arm wants to put your Starbucks barista out of business. Retrieved February 11, 2019 from <https://www.cnbc.com/2018/05/08/this-25000-robot-wants-to-put-your-starbucks-barista-out-of-business.html>.
- Hoffman, D. L., & Novak, T. P. (2018). Consumer and object experience in the internet of things: An assemblage theory approach. *Journal of Consumer Research*, 44(6), 1178–1204.
- Huang, M. H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155–172.
- Hullinger, J. (2016). What the Lowe’s robot will do for you – and the future of retail. Retrieved February 11, 2019 from <http://campfire-capital.com/retail-innovation/sales-channel-innovation/what-the-lowes-robot-will-do-for-you-and-the-future-of-retail/>.
- Kaplan, A., & Haenlein, M. (2019). Siri, Siri, in my hand: Who’s the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15–25.
- Kidd, C. D., & Breazeal, C. (2008, September). Robots at home: Understanding long-term human-robot interaction. In *Proceedings of the IEEE International Conference of Intelligent Robot Systems* (pp. 3230–3235).
- Kim, T., & Duhachek, A. (2018). The impact of artificial agents on persuasion: A construal level account. *ACR Asia-Pacific Advances*.
- Knight, W. (2017). Forget killer robots – bias is the real AI danger. Retrieved February 11, 2019 from <https://www.technologyreview.com/s/608986/forget-killer-robots-bias-is-the-real-ai-danger/>.
- Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11), 1238–1274.
- Kwak, S., Kim, Y., Kim, E., Shin, C., & Cho, K. (2013). What makes people empathize with an emotional robot? The impact of agency and physical embodiment on human empathy for a robot. In *Proceedings of the IEEE International Symposium Robot Human Interaction Community* (pp. 180–185).
- Lammer, L., Huber, A., Weiss, A., Vincze, M. (2014). Mutual care: How older adults react when they should help their care robot. In *Proceedings of the International Symposium of Social Study AI Simulation Behavior* (pp. 3–4).
- Larson, K. (2019). Data privacy and AI ethics stepped to the fore in 2018. Retrieved February 11 from <https://medium.com/@Smalltofeds/data-privacy-and-ai-ethics-stepped-to-the-fore-in-2018-4e0207f28210>.
- Lashinsky, A. (2019). Artificial intelligence: Separating the hype from the reality. Retrieved February 11, 2019 from <http://fortune.com/2019/01/22/artificial-intelligence-ai-reality/>.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Lee, A. Y., Keller, P. A., & Sternthal, B. (2009). Value from regulatory construal fit: The persuasive impact of fit between consumer goals and message concreteness. *Journal of Consumer Research*, 36(5), 735–747.
- Leung, E., Paolacci, G., & Puntoni, S. (2018). Man versus machine: Resisting automation in identity-based consumer behavior. *Journal of Marketing Research*, 55(6), 818–831.
- Levin, S. (2017). New AI can guess whether you’re gay or straight from a photograph. Retrieved February 11, 2019 from <https://www.theguardian.com/technology/2017/sep/07/new-artificial-intelligence-can-tell-whether-youre-gay-or-straight-from-a-photograph>.

- Li, D., Rau, P. P., & Li, Y. (2010). A cross-cultural study: Effect of robot appearance and task. *International Journal of Social Robotics*, 2(2), 175–186.
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to Medical Artificial Intelligence. *Journal of Consumer Research* (forthcoming).
- Lowy, J. (2016). Self-driving cars are ‘absolutely not’ ready for deployment. Retrieved February 11, 2019 from <https://www.pbs.org/newshour/science/self-driving-cars-are-absolutely-not-ready-for-deployment>.
- Luo, X., Tong, S., Fang, Z. & Zhe, Q. (2019). Machines vs. humans: The impact of chatbot disclosure on consumer purchases. Unpublished working paper.
- Mandel, N., Rucker, D. D., Levav, J., & Galinsky, A. D. (2017). The compensatory consumer behavior model: How self-discrepancies drive consumer behavior. *Journal of Consumer Psychology*, 27(1), 133–146.
- Marr, B. (2019). How robots, IoT and artificial intelligence are changing how humans have sex. Retrieved June 12, 2019 from <https://www.forbes.com/sites/bernardmarr/2019/04/01/how-robots-iot-and-artificial-intelligence-are-changing-how-humans-have-sex/#3679d398329c>.
- Martin, K. D., & Murphy, P. E. (2017). The role of data privacy in marketing. *Journal of the Academy of Marketing Science*, 45(2), 135–155.
- Martin, K. D., Borah, A., & Palmatier, R. W. (2017). Data privacy: Effects on customer and firm performance. *Journal of Marketing*, 81(1), 36–58.
- Mehta, N., Detroja, P., & Agashe, A. (2018). Amazon changes prices on its products about every 10 minutes — here’s how and why they do it. Retrieved February 11, 2019 from <https://www.businessinsider.com/amazon-price-changes-2018-8?international=true&r=US&IR=T>.
- Mende, M., Scott, M. L., van Doorn, J., Grewal, D., & Shanks, I. (2019). Service robots rising: How humanoid robots influence service experiences and food consumption. *Journal of Marketing Research*, 56(4), 535–556.
- Metz, C. (2018). Mark Zuckerberg, Elon musk and the feud over killer robots. Retrieved February 11, 2019 from <https://www.nytimes.com/2018/06/09/technology/elon-musk-mark-zuckerberg-artificial-intelligence.html>.
- Milgram, P., Takemura, H., Utsumi, A., & Kishino, F. (1995, December). Augmented reality: A class of displays on the reality-virtuality continuum. *Telematic and telepresence technologies*, 2351, 282–293.
- Miller, G. (2016). Bots, chatbots and artificial intelligence, what the evolution is all (aBot) about. Retrieved June 12, 2019 from <https://chatbotnewsdaily.com/bots-chatbots-and-artificial-intelligence-what-the-evolution-is-all-abot-about-a7e148dd067d>.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
- Moon, Y. (2003). Don’t blame the computer: When self-disclosure moderates the self-serving bias. *Journal of Consumer Psychology*, 13(1–2), 125–137.
- Mori, M. (1970). The Uncanny Valley. *Energy*, 7(4), 33–35.
- Motyka, S., Grewal, D., Puccinelli, N. M., Roggeveen, A. L., Avnet, T., Daryanto, A., et al. (2014). Regulatory fit: A meta-analytic synthesis. *Journal of Consumer Psychology*, 24(3), 394–410.
- Müller, V. C., & Bostrom, N. (2016). Future progress in artificial intelligence: A survey of expert opinion. In *Fundamental issues of artificial intelligence* (pp. 555–572). Springer, Cham.
- Parekh, J. (2018). Why Programmatic provides a better digital marketing landscape. Retrieved February 13, 2019 from <https://www.adweek.com/programmatic/why-programmatic-provides-a-better-digital-marketing-landscape/>.
- Pedersen, I., Reid, S., & Aspevig, K. (2018). Developing social robots for aging populations: A literature review of recent academic sources. *Sociology Compass*, 12(6).
- Power, B. (2017). How AI is streamlining marketing and sales. Harvard Business Review, June 12. Retrieved February 11, 2019 from <https://hbr.org/2017/06/how-ai-is-streamlining-marketing-and-sales>.
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J. F., Breazeal, C., et al. (2019). Machine behavior. *Nature*, 568(7753), 477–486.
- Reese, B. (2018). *The fourth age: Smart robots, conscious computers and the future of humanity*. New York, NY: Atria Books.
- Robinson, M. (2017). The tech giants of Silicon Valley are starting to rely on crime-fighting robots for security. Retrieved February 11, 2019 from <https://www.businessinsider.com/knightscope-security-robots-microsoft-uber-2017-5>.
- Sandoval, E. B., Brandstetter, J., Obaid, M., & Bartneck, C. (2016). Reciprocity in human-robot interaction: A quantitative approach through the prisoner’s dilemma and the ultimatum game. *International Journal of Social Robotics*, 8(2), 303–317.
- Schatsky, D., Katyal, V., Iyengar, S., & Chauhan, R. (2019). Why enterprises shouldn’t wait for AI regulation. Retrieved July 11, 2019 from <https://www2.deloitte.com/insights/us/en/focus/signals-for-strategists/ethical-artificial-intelligence.html>
- Schrift, R. Y., Wertenbroch, K., André, Q., & Frank, D.H. (2017). Threatening free will. Presentation at the *Symposium on Alienation and Meaning in Production and Consumption*, Technische Universität München.
- Shankar, V. (2018). How artificial intelligence (AI) is reshaping retailing. *Journal of Retailing*, 94(4), vi–xi.
- Sherwood, H. (2017). Robot priest unveiled in Germany to mark 500 years since reformation. *The Guardian*, may 20. Retrieved February 11, 2019 from <https://www.theguardian.com/technology/2017/may/30/robot-priest-blessu-2-germany-reformation-exhibition>.
- Shum, P., & Lin, G. (2007). A world class new product development best practices model. *International Journal of Production Research*, 45(7), 1609–1629.
- Şimşek, Ö. F., & Yalınçetin, B. (2010). I feel unique, therefore I am: The development and preliminary validation of the personal sense of uniqueness (PSU) scale. *Personality and Individual Differences*, 49(6), 576–581.
- Singh, J., Flaherty, K., Sohi, R. S., Deeter-Schmelz, D., Habel, J., Le Meunier-FitzHugh, K., et al. (2019). Sales profession and professionals in the age of digitization and artificial intelligence technologies: Concepts, priorities, and questions. *Journal of Personal Selling & Sales Management*, 1–21.
- Smith, C. (2019) Understanding China’s Confucian edge in the global AI race, *Quillette*, Feb. 14, 2019, Retrieved from <https://quillette.com/2019/02/14/understanding-chinas-confucian-edge-in-the-global-ai-race/> on July 6, 2019.
- Steels, L., & Brooks, R. (2018). *The artificial life route to artificial intelligence: Building embodied, situated agents*. Routledge.
- Syam, N., & Sharma, A. (2018). Waiting for a sales renaissance in the fourth industrial revolution: Machine learning and artificial intelligence in sales research and practice. *Industrial Marketing Management*, 69, 135–146.
- Tucker, C. (2018). Privacy, algorithms, and artificial intelligence. In *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- van Lente, H., Spitters, C., & Peine, A. (2013). Comparing technological hype cycles: Towards a theory. *Technological Forecasting and Social Change*, 80(8), 1615–1628.
- Verhoef, P. C., Stephen, A. T., Kannan, P. K., Luo, X., Abhishek, V., Andrews, M., et al. (2017). Consumer connectivity in a complex,

- technology-enabled, and mobile-oriented world with smart products. *Journal of Interactive Marketing*, 40, 1–8.
- Villasenor, J. (2019). Artificial intelligence and bias: Four key challenges. Retrieved February 11, 2019 from <https://www.brookings.edu/blog/techtank/2019/01/03/artificial-intelligence-and-bias-four-key-challenges/>.
- Wainer, J., Feil-Seifer, D. J., Shell, D. A., & Mataric, M. J. (2006, September). The role of physical embodiment in human-robot interaction. In *Proceedings of the IEEE International Symposium of Robot Human Interaction Communication* (pp.6-8).
- Wang, Y., & Kosinski, M. (2018). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of Personality and Social Psychology*, 114(2), 246–257.
- Weissman, J. (2018). Amazon created a hiring tool using a.I. it immediately started discriminating against women. October 10. Retrieved February 11, 2019 from <https://slate.com/business/2018/10/amazon-artificial-intelligence-hiring-discrimination-women.html>.
- Wilson, S. (2018). Big privacy: The data privacy compact for the era of big data and AI. Retrieved February 11, 2019 from <https://www.zdnet.com/article/big-privacy-the-data-privacy-compact-for-the-era-of-big-data-and-ai/>.
- Wilson J., Daugherty P., & Shukla P. (2016). How one clothing company blends AI and human expertise. *Harvard Business Review*, November 21. Retrieved February 11 from <https://hbr.org/2016/11/how-one-clothing-company-blends-ai-and-human-expertise>.
- Wirtz, J., Patterson, P. G., Kunz, W. H., Gruber, T., Lu, V. N., Paluch, S., & Martins, A. (2018). Brave new world: Service robots in the front-line. *Journal of Service Management*, 29(5), 907–931.
- You, Q., Jin, H., Wang, Z., Fang, C., & Luo, J. (2016). Image captioning with semantic attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4651-4659).
- Zhao, M., Hoeffler, S., & Dahl, D. W. (2012). Imagination difficulty and new product evaluation. *Journal of Product Innovation Management*, 29(S1), 76–90.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.