

Catalyst

Innovation and Design in the Age of Artificial Intelligence

Roberto Verganti, Luca Vendraminelli, and Marco Iansiti

At the heart of any innovation process lies a fundamental practice: the way people create ideas and solve problems. This “decision making” side of innovation is what scholars and practitioners refer to as “design.” Decisions in innovation processes have so far been taken by humans. What happens when they can be substituted by machines? Artificial Intelligence (AI) brings data and algorithms to the core of the innovation processes. What are the implications of this diffusion of AI for our understanding of design and innovation? Is AI just another digital technology that, akin to many others, will not significantly question what we know about design? Or will it create transformations in design that current theoretical frameworks cannot capture?

This paper proposes a framework for understanding the design and innovation in the age of AI. We discuss the implications for design and innovation theory. Specifically, we observe that, as creative problem-solving is significantly conducted by algorithms, human design increasingly becomes an activity of sensemaking, that is, understanding which problems should or could be addressed. This shift in focus calls for the new theories and brings design closer to leadership, which is, inherently, an activity of sensemaking.

Our insights are derived from and illustrated with two cases at the frontier of AI—Netflix and Airbnb (complemented with analyses of Microsoft and Tesla)—which point to two directions for the evolution of design and innovation in firms. First, AI enables an organization to overcome many past limitations of human-intensive design processes, by improving the scalability of the process, broadening its scope across traditional boundaries, and enhancing its ability to learn and adapt on the fly. Second, and maybe more surprising, while removing these limitations, AI also appears to deeply enact several popular design principles. AI thus reinforces the principles of Design Thinking, namely: being people-centered, abductive, and iterative. In fact, AI enables the creation of solutions that are more highly user centered than human-based approaches (i.e., to an extreme level of granularity, designed for every single person); that are potentially more creative; and that are continuously updated through learning iterations across the entire product life cycle.

In sum, while AI does not undermine the basic principles of design, it profoundly changes the practice of design. Problem-solving tasks, traditionally carried out by designers, are now automated into learning loops that operate without limitations of volume and speed. The algorithms embedded in these loops think in a radically different way than a designer who handles the complex problems holistically with a systemic perspective. Algorithms instead handle complexity through very simple tasks, which are iterated continuously. This paper discusses the implications of these insights for design and innovation management scholars and practitioners.

Introduction

The adoption of artificial intelligence (AI) has received enormous attention across virtually every industrial setting, from healthcare delivery to automobile manufacturing. In combination with the ubiquity of digital sensors, networks, and software-based automation, AI is transforming our economy and

defining a new industrialization age. From Alibaba to Airbnb, this “Age of AI” is defined by the emergence of a new kind of firm, based on a digital operating model, creating unprecedented opportunities and challenges (Iansiti and Lakhani, 2020a, 2020b, 2020c).

As firms evolve to embrace an increasingly AI-centric operating model, they are digitizing a growing number of important business processes, removing human labor and management from the execution of many critical operating activities. For example, unlike processes in traditional firms, no worker sets the

Address correspondence to: Roberto Verganti, Stockholm School of Economics, Hölländargatan 32, Box 6501, SE-113 83 Stockholm, Sweden. E-mail: roberto.verganti@hhs.se.

price on an Amazon product or qualifies a business for a loan at Ant Financial. While humans develop the algorithms and write the software code, the actual real-time creation of the solution is automated and enabled entirely by digital technology.

As the economy continues to transform, innovation processes are also changing rapidly, making use of sensors, digital networks, and algorithms. Whether the product consists entirely of software, as with an iPhone app, or whether it is a more traditional hardware-centric artifact, as in a Tesla automobile, modern products are increasingly connected to the organization that created them, providing a continuous flow of data that details many aspects of the user experience. In addition, the software embedded in the products themselves enables information flowing the other way, from the firm to the user, enabling a specific solution for a specific person, constantly improving the experience in real time. These instant two-way interactions characterize an increasing range of goods and services,

from Netflix video streaming to a Tesla Model 3. Effectively, these innovative solutions evolve in real time as the user experiences them.

It is important to note that to bring about the kinds of dramatic changes we are describing, we do not need a particularly advanced notion of AI. AI need not be indistinguishable from the human behavior, or capable of simulating human reasoning—what is sometimes defined as “strong AI” in the field of computer science. We do not need a perfect human replica to prioritize content on a social network, optimize the recipe for a perfect cappuccino, analyze customer behavior patterns, understand the implications of design trade-offs, or personalize a product. We need only a computer system to perform simple tasks that were traditionally performed by human beings, such as recognizing images, or processing natural language. This is what traditionally defines “weak AI” (Iansiti and Lakhani, 2020a). Imperfect, weak AI, typically powered by the exploding field of machine learning, is already enough to create significant change when replicated at scale.

AI, as defined above, profoundly transforms the context where innovation takes place. Why? AI is inherently a decision-making technology: it offers opportunities to automate many tasks relating to learning and devising solutions. When AI is applied to the context of innovation, it may therefore transform how decisions in innovation are made, especially in relation to how novel solutions (whether a good, a service, or a process) are created and tested. This decision-making practice at the heart of innovation is what scholars refer to as design (Liedtka, 2015). Indeed, ultimately, to design is to “devise courses of action aimed at changing existing situations into preferred ones” (Simon, 1982, p. 129). Investigating how AI affects the innovation processes therefore requires an exploration of how it affects the design.

In this paper we explore implications of AI for the design and innovation management by exploring the strategies of pioneering organizations such as Netflix and Airbnb. Our analysis addresses three sets of questions:

- Questions about AI and the *practice* of design: To what extent is AI likely to change the way design is practiced, that is, *which* decisions are made and *how*? Is the transformation of the context induced by AI changing the design *process* and the *objects* of the design actions? For example, which decisions can be automated and which ones cannot?

BIOGRAPHICAL SKETCHES

Roberto Verganti is professor of leadership and innovation at the House of Innovation of the Stockholm School of Economics, where he co-directs The Garden—the Center for Design and Leadership. He is also the co-founder of Leadin’Lab, the laboratory on the Leadership, Design and Innovation of Politecnico di Milano. Roberto is the author of *Overcrowded. Designing Meaningful Products in a World Awash with Ideas*, published by MIT Press in 2017, and of *Design-Driven Innovation*, published by Harvard Business Press in 2009, which has been selected by the Academy of Management for the George R. Terry Book Award as one of the best six management books published in 2008 and 2009. Roberto serves on the Advisory Board of the European Innovation Council of the European Commission. (www.verganti.com)

Luca Vendraminelli is a PhD candidate at the University of Padova, Italy and a visiting fellow at LISH, the Laboratory for Innovation Science at Harvard University. His interests revolve around the mechanisms of digital transformations in traditional companies. He is investigating how to govern digital transformation by means of strategy, and the effects of Artificial Intelligence on operating models.

Marco Iansiti is the David Sarnoff Professor of Business Administration at Harvard Business School, where he is also a co-director of LISH, the Laboratory for Innovation Science, and of the Digital Initiative at HBS. His work has appeared in scientific journals such as *Journal of Product Innovation Management*, *Strategic Management Journal*, *Management Science*, *Organization Science*, *Research Policy*, and *Industrial and Corporate Change*. More recently, his writings on innovation and digital transformation made HBR’s list of top 10 papers of the year for three of the last four years. Iansiti has authored or coauthored several books, including *Technology Integration*, and *The Keystone Advantage*, both published by Harvard Business School Press. His latest book, co-authored with Karim R. Lakhani, is *Competing in the Age of Artificial Intelligence* published by Harvard Business School Press (2020).

- Questions about AI and the *principles* of design: If AI induces significant changes in design practice, are these changes putting the fundamentals of design in question? Is, for example, user centeredness, questioned? Is design practice, in the age of AI, informed by significantly different principles?
- Finally, questions about AI and the *theory* of design and innovation: What are the implications for the theoretical frameworks that we use to interpret design and innovation? Does the widespread adoption of AI call for new research questions and for a new understanding of how design drives innovation in organizations?

The article is structured as follows. We start by introducing the principles of design, with a special focus on recent developments in Design Thinking. Then, we introduce a framework that enables to compare the traditional human-intensive design practice with design practice in the age of AI. The framework is then illustrated with the cases of Netflix and Airbnb. Next, we discuss the cases (with support of additional information from the experiences in Tesla and Microsoft) to analyze the extent to which the design principles and practice are affected by AI. We then conclude with an analysis of implications for design and innovation theories and scholars.

Design and its Operating Context

To investigate whether and to what extent AI transforms our understanding of design, we frame our discussion according to two levels of analysis (Orlikowski, 2010): practice and principles. Design *practice* refers to the phenomenology of design in a specific context: its *process* (“how” design decisions are made; through which phases, methods, tools, or collaborative practices) and the *object* of design (which design decisions are made; which novel solution it creates, whether a good, service, or process). Design *principles*, instead, refer to the perspective and philosophy that inform the act of designing, and that constitute an ontology of what design is. The distinction between these two levels of analysis enables us to better discern how AI might affect the design. Is AI changing the way we design, or is it acting at a deeper level by reframing the basic principles that inspire the act of designing? To answer this question, we start by introducing the principles of design, as they emerge from the current discussion

on design and innovation theory. We then illustrate how these principles have been instantiated into design practice before the advent of AI. Finally, we introduce a framework to analyze how these principles are enacted in the context of AI.

The Principles of Design

What are the principles that inform the practice of design? The scientific debate on the ontology of design has developed in the realm of design theories, with a rich set of contributions (see for example Galle, 2002; Love, 2000; Margolin, 1989; Margolin and Buchanan, 1996). Given our focus on design practices in organizational contexts, we take a more specific perspective: Design Thinking. This perspective leverages the body of design theory literature and adapts it to interpret how design-driven innovation can happen in the context of business. Although the term Design Thinking suffers from some ambiguity, the efforts of management scholars to distill its principles converge toward the three essential factors (see Calabretta and Kleinsmann, 2017; Dell’Era, Cautela, Magistretti, Verganti, and Zurlo, 2020; Micheli, Wilner, Bhatti, Mura, and Beverland, 2019; Seidel and Fixson, 2013; and especially Liedtka, 2015, for a re-composition of the principles of design thinking with the principles of design theories):

- *People-centered*: innovation, when driven by design, is inspired by empathy with users. Rather than being driven by the advancements of technology and by what is possible, design-driven innovation stems from understanding a problem from the user perspective, and from making predictions about what could be meaningful to her. For example, we can recognize this principle in the practice of ethnographic research.
- *Abductive*: design adopts a creative approach to solve problems, which sets it apart from other problem-solving practices in management, as clarified by Boland and Collopy (2004): “We portray the manager as facing a set of alternatives from which a choice must be made. This decision attitude assumes it is easy to come up with alternatives to consider, but difficult to choose among them. The design attitude toward problem-solving, in contrast, assumes that it is difficult to design a good alternative, but once you have developed a



Figure 1. Design Practice in the Context of Traditional Human-Intensive Operating Models [Colour figure can be viewed at wileyonlinelibrary.com]

truly good one, the decision about which alternative to select is trivial” (p. ix). Design therefore implies to *imagine* the new rather than *finding* a solution within a set; as Simon (1982) states, design is “concerned not with the necessary but with the contingent—not with how things are but with how they might be” (p. xii). From the perspective of logical inference, this implies that design solves problems through abductions: rather than leveraging solely deductive reasoning (how things are) and inductive reasoning (how things likely are), design creates through abductive reasoning (by making hypotheses about how things might be). This is why design is often associated with creativity and ideation rather than analysis. For example, we can recognize this principle in the practice of brainstorming.

- *Iterative*: abductions are continuously adapted and improved through fast testing cycles. The prototypes that are built in these cycles act as a “playground” for conversation and learning (Schrage, 1999). They engage the team and users in iterations in which solutions are tested and refined, until a satisfactory result is achieved. For example, we can recognize this principle in the practice of building rudimentary mockups early in the design process.

Design in the Context of Traditional Operating Models

The way design principles are enacted into practice depends of course on the operating context in which design takes place. Most design practices we know today rely on human decision-making. Because of

this labor-intensive design context, it is not practically possible nor economically convenient to design a different solution for each individual user. Products (goods and services) are therefore designed for segments of users (see the phase “design” in Figure 1). Then, products are manufactured at scale, through complex production architectures which include possibilities for customization (“make”). Finally, they are delivered for “use” (see also Clark and Fujimoto, 1991).

After a product is released, the context evolves. For example, the market changes, or new technological opportunities emerge. In addition, organizations can learn new insights from how customers actually use the existing product. Yet, as the operating model entails significant effort and investment to redesign a product, innovation is postponed until the marginal value of a new product supersedes the cost of its design. At this point, a new design cycle starts.

A structure of this kind therefore implies a significant separation in time between two consequent design initiatives. During product use, learning cycles are frozen and, consequently, solutions become

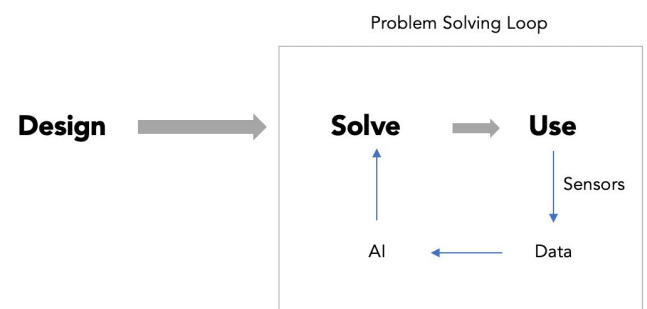


Figure 2. Design Practice in the Context of AI Factories [Colour figure can be viewed at wileyonlinelibrary.com]

rapidly “old.” New learning and ideas may only be incorporated in future solutions released in lumps, episodically and for customer segments.

Design in the Context of AI Factories

As discussed above, traditional design activities are human intensive. AI offers the opportunity to revolutionize this scenario. To understand how, we have explored cases of organizations that are pioneering the use of AI in design, namely Netflix and Airbnb. The observation of these organizations allowed us to develop an original framework (Figure 2) that describes how design practice can be articulated in the age of AI. Let’s briefly describe the main elements of the framework, before delving into the description of these illustrative cases in the next section.

To design implies making a number of decisions. A few of them are highly sophisticated and conceptual. But most decisions, especially during development, are narrow and ask for specific problem-solving skills. Examples of these detailed decisions are the choice of the functional shape of an object, the details of a product interface, or which information to display on a screen. There are plenty of detailed problems to be addressed during design. AI offers the intelligence to solve them.

In the context of AI factories (i.e., organizations that make intensive use of AI in their operating models—see Iansiti and Lakhani, 2020a, 2020b) a specific solution, that is, what an individual user actually interacts with, is designed by an AI engine in what we call “problem-solving loops.” Loops collect real time data (insights) from customer interactions or from the ecosystem in which the firm lies. These data can immediately inform the AI embedded in the product, which has problem-solving capabilities (from recognizing objects to processing natural language, from making predictions to drawing conclusions). If properly conceived, an algorithm can autonomously generate a new specific solution for that precise user, with no human effort involved. Even more, as new data are continuously collected, and the AI engine embeds learning capabilities, the problem-solving loops improve their predictions about user needs and behaviors and therefore design better solutions over time.

In an AI-powered system, many development decisions are therefore made through problem-solving

loops that are autonomous and human capital-free. The work of humans is to conceive the foundations for a new offering and design these problem-solving loops (see the phase “design” in Figure 2). The loops will then replace people with technology in the development of a specific solution: they are easy to scale without redesign, and can provide a variety of solutions without large additional investments in R&D.

AI-Empowered Design in Practice

To see how the framework of Figure 2 works in practice, we examined the cases of Netflix and Airbnb. We selected these two cases, as, being at the frontier of AI applications, they offer a glimpse into the future of design in the context of an AI-centric firm.

Netflix and the Data-Driven, Design Thinking Machine

Netflix has completely transformed the media landscape by harnessing the power of big data and AI. The core of Netflix is its data and AI-centric operating model. It is powered by software infrastructure that gathers data and trains and executes algorithms that drive virtually every aspect of the business, from personalizing the user experience to picking winning movie concepts for its next productions. In this section, we detail the Netflix approach to design, by digging into some of the machine learning techniques that Netflix has deployed into its problem-solving loops.

Netflix started to leverage AI at least as early as 2010, to fuel its recommendation engine. In 2014, Netflix expanded its approach to invest extensively in understanding user behavior and develop a personalized streaming experience for each user. The application screens that a user sees today are “designed in real time” by a machine. Many boundaries and parameters are specified by human designers at the outset of the process. But the decisions about which movies to show, how to display them, which pictures to represent them with, and many other design decisions are done by algorithms embedded in the AI problem-solving loops. Let’s dig into these algorithms, which effectively resemble different aspects of a process of design.

The basic problems most AI systems try to solve to shape a design experience relate to predicting an

outcome. The tool for making that prediction is an algorithm—the set of rules a machine follows to solve a particular problem. AI can incorporate many types of algorithms (Domingos, 2012). Some of them have a built-in process for updating and improving, most often based on “Markov decision processes,” which seek to model a sequence of actions, each shaped by a policy, and followed by a reward. One example would be the Netflix algorithms that dynamically update its user interface, based on the actual behavior of the user, as indicated by her clicks (while the policy decides what is displayed, the click is the “reward”).

While applications have exploded over the last decade, the foundations of algorithm design have been around for a while. The conceptual and mathematical development of classical statistical models like linear regression, clustering, or Markov chains dates back more than a hundred years. Today’s neural networks were initially developed in the 1960s and are only now being put to use at a scale with production-ready outputs. The vast majority of production-ready and operational AI systems at Netflix use one of three general approaches to developing accurate predictions using statistical models, also known as machine learning. These are supervised learning, unsupervised learning, and reinforcement learning.

Supervised learning. The basic goal of supervised machine-learning algorithms is to come as close as possible to an expert (or an accepted source of truth) in predicting an outcome. The classic case is analyzing a picture and predicting whether the subject is a cat or a dog. In this case, the expert would be any human being with good eyesight who could label photos as cat or dog. The first step in supervised learning is to create (or acquire) a labeled data set. The data are then split between training and validation. As we compare the algorithmic model’s prediction of the outcome to the validated labeled outcomes, we can determine if we are satisfied with the error between prediction and expert. If we are not satisfied, we can go back and choose a different statistical approach, get more data, or work on identifying other features that may be helpful in making a more accurate prediction. Netflix uses supervised learning in a variety of scenarios. For recommendations, Netflix has used labeled data sets made up of actions and results (e.g., movies chosen and liked) by people who are deemed by the algorithm to be similar to a given user. A large data set of user choices calibrated by characteristics of the

user and of the decision context can lead to effective recommendations.

Note that supervised learning resembles elements of human design, as instantiated earlier in the first principle of design (people centered). Just as human designers immerse themselves in the context of use and observe all possible aspects of the user experience, the algorithms are trained by a relevant stream of *user data*, with significant information on the *context of use* (e.g., the type of device, time and place of action, and so on). The richer the stream of data, the more the problem-solving loops are user centered.

Unsupervised learning. Unlike supervised learning models, which train a system to recognize known outcomes, the primary application of unsupervised learning algorithms is to discover insights in data with few preconceptions or assumptions. Whereas in supervised learning the data inputs are labeled with a given outcome, unsupervised learning algorithms aim to find “natural” groupings in the data, without labels, and uncover structure that may not be obvious to the observer. In our example of photos of cats and dogs, an unsupervised learning algorithm might find several types of groupings. Depending on how the clusters are structured, these could end up separating cats and dogs, or indoor and outdoor photographs, or pictures taken during day or night, or virtually anything else. In these cases, one does not know exactly what to look for, but is searching for related groups. Netflix uses unsupervised learning to discover related groups of customers or to create different versions of the user interface that match different usage patterns. Even more advanced, Netflix uses data and AI algorithms to predict which content to create in the first place. The first application of predictive analytics was back in 2013, to evaluate the potential of *House of Cards*, in collaboration with Media Rights Capital. The new series was a hit and Netflix continued to develop content in response to detailed predictive analytics on market and user behavior. Cindy Holland, vice president of original content, noted in an interview: “We have projection models that help us understand, for a given idea or area, how large an audience size might be, given certain attributes about it. We have a construct for genres that basically gives us areas where we have a bunch of programs and others that are areas of opportunity” (Spangler, 2018).

Note that unsupervised learning is a relatively unstructured design process, where the patterns that

emerge at the end do so based strictly on the observations (the data) and are not set up at the outset of the process. Although at its core, the algorithm simulates induction, when perpetuated on an extremely large quantity of data unsupervised learning provides insights and hypotheses that mirror the abductions of humans, or the dynamics of ideation and brainstorming. Hence, unsupervised learning also embeds the basic perspectives of design thinking into the problem-solving loops of the AI factory.

Reinforcement learning. Reinforcement learning makes up the third machine learning paradigm and is the closest in structure to a traditional design process. The applications of reinforcement learning may be even more impactful than those of supervised and unsupervised learning. Rather than starting with data on an expert's view of the outcome, as in supervised learning, or with a pattern and anomaly recognition system, as in unsupervised learning, reinforcement learning just requires a starting point and a performance function. The system starts somewhere and probes the space around the starting point, using as a guide whether it has improved or worsened the performance of the algorithm. The key trade-off is whether to spend more time *exploring* the contextual complexity beyond the current understanding or *exploiting* the model built so far to drive decisions and actions.

Let's say we take a cable car up a tall mountain and we want to walk our way down. It is a really foggy day and the mountain does not have any clearly marked paths. Since we cannot just see the best way down, we have to walk around and explore different options. There is a natural trade-off between the time we spend walking around and getting a feel for the mountain, and the time we spend actually walking down once we believe we have found the best path. This is the trade-off between exploration and exploitation. The more time we spend exploring, the more we will be convinced we have the best way down, but if we spend too long exploring, we will have less time to exploit the information and actually walk down.

This is pretty close to the way the Netflix algorithm actually personalizes movie recommendations and the visuals they are associated with. Through the analysis of user data, Netflix recognized that viewers have enormous diversity in taste and preferences. So, the Netflix team decided that each user should be shown a cover artwork specifically designed for her, drawn

from the frames of a movie. The artwork would highlight the aspects of the title that are specifically relevant to that specific user (Chandrashekar, Amat, Basilico, and Jebara, 2017). The problem was complicated, as the Netflix team needed to figure out which movie selection to present, and then, which artwork to combine with that movie to maximize the match between user and recommendation. A single season of an average TV show (about 10 episodes) contains nearly 9 million total frames. Asking creative editors to efficiently sift through that many frames of videos to design an artwork that would capture the audience's attention would be tedious and ineffective. Designing an artwork for each specific user according to his or her preferences would simply be impossible. But an AI factory, and in particular reinforcement learning loops, can address this design problem effectively. In a way similar to our previous example (finding our way down the mountain), Netflix uses reinforcement learning (and in particular multi-arm bandit algorithms) to spend some time exploring options, and some time exploiting the solution offered by its models. To explore visual options and refine the prediction model, Netflix systematically randomizes the visuals shown to a user. Netflix then exploits the improved model to show a specific user a slew of recommendations with improved visuals. The Netflix service continues to improve dynamically, by automatically cycling between periods of exploration and exploitation, designed to learn the most about the preferences of a complex human being, and maximize the engagement of this specific user over the long haul.

Note that with its emphasis on balancing exploitation and exploration, reinforcement learning resembles the process of human design in many facets, and in particular the principle of iterations enunciated earlier. Just as we see with traditional design approaches, opening the funnel with broad exploration can lead to more interesting and innovative decisions, but must be balanced with the increased challenge in making sure the exploitation phase converges on a usable solution.

In its earliest days two decades ago, the Netflix operating model consisted of shipping DVDs. With this mail delivery service, Netflix could only track which titles users viewed, how long they kept a DVD, and how they rated each title, but they could not monitor actual viewing behavior. Although Netflix already recognized the importance of using data to improve customer experience, the heaviness of its assets and operations

gave limitations to its capability to design. But when in 2007 Netflix launched its streaming service, the company seized the opportunity to transform its operating model into an AI enabled one. With streaming, Netflix could track the full user experience—when viewers pause, rewind, or skip during a show, for example, or what device they watch it on. This enabled to design several problem-solving loops that bring design principles to its extreme level: a different solution for every single user, designed and delivered on the fly. As Joris Evers, Netflix's then-chief of communications, says “there are 33 million different versions of Netflix” (Carr, 2013).

How Airbnb Reframed the Design Practice in the Hospitality Industry

The case of Netflix offered an opportunity to understand the design practices of an organization, as it transitions from a traditional operating model to an AI enabled one. In particular, it illustrated how problem-solving loops work, the different configurations they can take, and how they enable to create people-centered solutions. How does this new form of design compare to other practices? To this purpose, the hospitality industry offers interesting insights, as its competitive arena contrasts players with traditional and AI-powered operating models.

Achieving people centricity in hospitality businesses is extremely complex by nature, because the context is characterized by diversity in many things, including cultures, ages, backgrounds, and travel purposes. As an indication of cultural complexity, consider that the chatbot of Booking.com translates 43 different languages. In the face of this kind of complexity, the traditional operating model of the industry was based on heavy investments in real estate (hotels and their spaces and rooms) and labor-intensive processes, with people that need to be hired, educated, and coordinated. Rooms of asset-heavy companies therefore need to be designed in a more or less standardized fashion, and they remain static for a significant span of time. Similarly, the user experience, and the related back end of the service, is designed and formalized to ensure that quality standards are always respected. This kind of traditional operating model creates significant challenges in delivering an experience that can fit the individual users.

To address the challenge, the industry has witnessed in the last decades some of the most frequent and

popularized initiatives of consolidated design thinking initiatives. Examples of design thinking applications in the industry included projects conducted by IDEO for Intercontinental Hotels Group. One project, for example, targeted short-stay travelers and aimed to create a convenient experience. Another targeted business travelers and resulted in the design of proper spaces for meeting and working. Yet another project focused on revamping the Holiday Inn Express brand by redesigning everything from how you check in to the look and feel of the room itself (Wilson, 2015). Each of these projects was framed according to a linear design practice, typical of traditional operating models: running ethnography to understand stakeholders' needs, ideating to formulate effective experiences for the target segment, and relying on rough prototyping for the identified solutions. To this purpose the innovation trend within the sector has been to create innovation labs that are spacious places where design teams can prototype rooms on a 1:1 scale. Sometimes, innovation in room design was taken to the extreme, and it was conducted directly on site: Marriott and Hilton selected real hotels to run beta-tests, where customers could directly get in touch with new ideas. The projects were then frozen into a design (of rooms, processes, or IT applications) that the asset-heavy operator could deliver in a proper consistent way at scale.

In the last part of the 2000s, a significant transformation of the hospitality sector began. New companies with a lighter operating model entered the industry, colliding with traditional asset-heavy business models (Iansiti and Lakhani, 2020c). Airbnb addressed a similar need to Marriott: providing space to guests who needs it. Consigning the onus of managing the operations to the hosts, Airbnb was able to get over traditional growth bottlenecks, such as the necessity to acquire rooms in order to scale. The capability of Airbnb to offer a state-of-the-art solution for each individual user depends on two factors.

First is the breadth of design options. For example, in 2017, Airbnb (which was founded in 2007), spread over more than 190 countries and 80,000 cities, and counted more than three million hosts: three times Marriott International's rooms, although it was founded in 1927. And, even more, these three million rooms were all different from the other three million designs. Traditional innovation practices in asset-intensive businesses could not create such a variety of physical designs.

Second, this enormous breadth of options had to be connected to the needs of each individual user. And here is where the AI comes into place. Airbnb collects an enormous amount of data from the interaction with each user. Since 2016, the data science team has developed an extensive logging within the booking flow that allows them to collect insights on what guests see, how they react to different types of interfaces, how much time they spend on a listing page, how long they take to make a booking request, or the exact time in which they decide to go back to search (Dai, 2017). When a customer interacts with Airbnb's search engine, a new event log (i.e., a list of user activity event data) is sent to a central repository. These logs pile up and detail the customer profile, with her preferences and behavior (Mayfield, Puttaswamy, Jagadish, and Long, 2016). Every time a customer reconnects to the service in search for a new traveling experience, Airbnb replies by instantaneously closing his problem-solving loop: data are extracted from the repository and processed by an AI-engine to create a new solution, personalized not only for the customer herself, but for the specific interaction.

The system works similarly to the case of Netflix. Even more interesting, Airbnb is a "two-sided platform"; that is, it interacts in real time with two categories of users (guests on one side, and hosts on the other). Its AI factory has therefore different problem-solving loops that work in parallel for each specific user type. An example of a problem-solving loop that provides effective solutions to hosts is how Airbnb designs the price of each individual list in an instantaneous and dynamic way: by ticking a box, a host accepts Airbnb's AI-engine to leverage data streams to automatically refine the price of their accommodation, within a price range. The AI-engine processes a vast amount of information collected from the ecosystem (Chang, 2017), such as the check-in lead time variation as the check-in date approaches, the listing popularity (i.e., how many people search around the host's area and how many of them click into the host's page), and the booking history, to understand how customers are reacting to price variations. The outcome is that the price is designed in the moment every time a guest asks for that specific property (Srinivasan, 2018).

AI offers the capability to perform different problem-solving loops independently on both sides of the platform, overcoming one of the limitations of traditional operational models: the need to balance

the requirements of different stakeholders. This design strategy enabled Airbnb to quickly become a central node of its network. The operating model was immensely scalable and allowed Airbnb to improve the quality of its services for both sides of the platform, thus enriching the communities of users and hosts at the same time. In a way, the case showcases the ultimate people centeredness: it provides solutions targeted to each individual person, across both users and hosts, driving personalization in a dynamic way that improves through constant iteration. It would have been impossible to achieve this with traditional design practices.

Discussion

As AI is diffusing in our society, scholars and practitioners wonder how this will impact our understanding of innovation and design. Our preceding exposition has important implications for design and innovation scholars and practitioners.

Artificial Intelligence and Design Practice

Up until today digital technologies have mainly spread into the operations of organizations, reducing the costs and time of manufacturing and delivering products and services. But the design of those products and services has largely remained a human intensive process. Referring back to Figure 1, even if the "making" was fast and cheap, the "designing" was heavy in time and resources. It was necessarily an intermittent activity, conducted in large projects and for a segment of users.

AI dramatically changes this scenario: it moves digital automation upstream, from manufacturing to design. Note that automation could be simply limited to accelerate traditional design tasks. For example, Airbnb is developing an AI system that can recognize sketches of customer experience hand-drawn by a designer on a drawing board and automatically render them into specifications for software engineers (Saarinen, 2017; Schleifer, 2017). If this were the only kind of use of AI, the essence of design practice would remain untouched: innovators would do what they did in the past (i.e., to draw components of the customer experience and translate them into specifications), but faster. However, Netflix and Airbnb go well beyond. They bring

automation directly into problem solving; that is, in the definition of detailed design choices: which interface to show to a specific user, which content to create, how to position a product compared to competitors. In this new context, designers and engineers do not simply make those decisions faster. They just do not make them, as they are delegated to AI. In other words, AI is the stimulus for an epiphany in the way we look at design (Magistretti, Dell’Era, and Verganti, 2020; Verganti, 2009, 2011a, 2011b). This has profound implications both in terms of the object and of the process of design.

The new object of design. The first dramatic change is in the object of design practice (the “what” of design). In human-intensive design, humans develop a product down to the level of details: for example, which image to be displayed on a screen. Conversely, with AI, the specific solution experienced by an individual user (i.e., what she actually sees on the screen of her mobile phone), is not only delivered but also *designed* by a problem-solving loop powered by AI. What humans do, in the context of AI, is not to design solutions (these are generated by the AI engine), but to design these problem-solving loops.

This change of object has disruptive implications. Especially because most AI algorithms do not reason like humans, that is, they do not just replicate and automate the thinking of an engineer or a designer; they work in a different way. Most of the applications we discussed in the cases of Netflix and Airbnb are instances of weak AI: they are focused on a combination of simple tasks (such as recognizing a shape in an image or if two images have different shapes) which are not nearly as sophisticated as the human thinking process they replace. Yet, by replicating these tasks millions of times (and by nurturing them with masses of data), weak AI can provide complex predictions, which even surpasses human capabilities.

The consequences are important. How do you design problem-solving loops? How do you conceive design rules that are based on extremely simple tasks, but that once replicated time and time again, can autonomously provide extremely complex solutions to users? Engineers and designers are not educated this way. Their mental frames are trained to systemically embrace complex tasks. To leverage the power of AI, they need an unprecedented capability: to imagine what a dumb system can do when operating at scale.

The new process of design. As the object of design changes (from designing solutions to designing problem-solving loops) the process of design (the “how” of design) changes as well. This is evident if we compare Figure 1 with Figure 2 previously illustrated: in the context of AI factories the design process is split into two chunks. First, a human-intensive design phase where the solution space is conceived and the problem-solving loops are designed; and then, an AI-powered phase, where the specific solution is developed for a specific user by the algorithm. As this second chunk of the process requires virtually zero cost and time, the development of the solution can be activated for each individual user, in the precise moment in which she asks for it. This in turn enables leveraging the latest available data and learning, and therefore creating, every time, a better novel solution. There are no more product or service blueprints that act as buffers between design and use. Design, delivery, and use—they all happen, in part, simultaneously.

Although this new practice is clearly visible in the realm of digital experiences based on software (such as Netflix and Airbnb), it is also gaining traction in industries based on physical products. Take for example the case of Tesla. Its operating model reflects those of Netflix or Airbnb, as it gathers massive amounts of data to design user experiences. However, to enable problem-solving loops Tesla is confronted by a tangible “hindrance”: the actual car. Hardware cannot be designed (yet) in real time, remotely and automatically. To unleash the power of AI, Tesla had therefore to reimagine the design of the car, acting in two diverse directions. First, it got rid of all the physical interacting elements (e.g., buttons) to embed most of the controls into digital user interfaces (e.g., into the large central touchscreen; Lambert, 2018). Second, it overloaded cars with sensors to collect data. Data are drawn from external sources (typically ultrasound equipment, GPS input, cameras, radar transmitters, and LIDAR) as well as internal ones. As cars go, sensors collect data and train Tesla’s learning algorithms. Interestingly enough, some of these sensors are “silent,” meaning they are not already used to provide direct value to customers, but placed “in perspective.” They are activated remotely after product release to enable new loops and provide new services to customers. The Model 3, for example, has been armed since 2017 with a cabin-facing camera placed in the rearview mirror. This camera was initially dormant (Lambert, 2017). Only in June 2019 was the camera used, thanks

to new software updates, to recognize occupants and adapt some of the hardware's adjustable components, such as the seats, vehicle mirrors, music, or driving mode preferences, in accordance with a specific user profile (Lambert, 2019).

Artificial Intelligence and Design Principles

Our cases show that in the context of an AI factory, design practice changes dramatically both in terms of the object and process of design. Does AI also undermine the core principles that underpin design? In other words, is this new design practice still people centered, abductive, and iterative? Or is it rooted in different principles? Our observations suggest that AI does not question the fundamental principles of design thinking. Rather, it further reinforces them.

To support this statement, we start from the findings of an extensive study of AI-powered strategies conducted by one of our coauthors (Iansiti and Lakhani, 2020a). The study shows that AI affects the operating model of an organization by eliminating three *limitations*: scale, scope, and learning. The cases discussed in this article show that AI removes these limitations also in innovation processes, empowering design's principles to be people centered, to create abductions, and to innovate through iterations.

Scale and people centeredness. Traditional design practice has significant scale limitations. Being one of the most intensive human-based activity, it requires the investment of significant resources and time. These scale limitations pose substantial constraints to people centeredness, as it is unreasonable to design a solution every time a user needs it. Products are instead designed for customer segments or average user archetypes (hence the use of “personas” in classic design thinking processes).

AI removes significant scale limitations in design, as the development of specific solutions is performed by machines. This enables the achievement of ultimate levels of people-centeredness. In fact, as seen in the case of Netflix, supervised learning leverages a rich stream of data on *each individual user*. This focus on individuals can be scaled with no limitations on the number of users and the complexity of data. As a consequence, the solution that a specific user experiences (e.g., what a user sees in the screen of the Netflix

application) has been developed just for her, *on the basis of her own data*. Interestingly, the relationship between scale and people-centeredness is now inverted. In human-intensive design, the larger the number of users and the complexity of insights, the more difficult it is to focus on individuals. In the context of AI factories, the larger the number of users and the richer and more complex the stream of data, the better the predictions of the machine on the behaviors of individuals. An even more advanced example is provided by Airbnb. Here the organization has to deal simultaneously with different categories of individuals: hosts and guests. Not only do the learning loops not suffer by this increase of complexity, but they also benefit from the integrated elaboration of data from both sides of the market.

Scope and abductions. Human-intensive design practices also have significant limitations in scope. Products are designed for a specific industry and with a specific target. Once they are released, they are unlikely to be applied in a different context. A car is designed to be a means of transportation. Moving from there to entertainment services is unlikely to happen. Limitations of scope are significant even within the same industry. Consider the case of Intercontinental Hotels Group, previously illustrated. The solutions developed by IDEO to address short-stay travelers and business travelers required different design initiatives, by different teams, and different brands of the same organization. The scope limitations of human-intensive design pose therefore significant constraints. Once a design brief is defined and frozen, creativity can happen only within the space of that brief.

AI enables the removal of many limitations in scope. In the context of AI factories, a design brief is fluid and can be reframed even after a product has been released. For example, we have seen how Netflix uses unsupervised learning to find new patterns in customer tastes that were not set up at the outset of the process. These predictions are used to support abductions in imagining new movie series. AI also makes it easier to imagine radically new services. Consider for example Airbnb, which has expanded into “travel experiences,” by offering guests the possibility to take a horse ride on a beach or hire musicians. To enter this new industry Airbnb leverages the same AI factory that powers the traditional hospitality service of AI.

Similarly, Tesla leverages the learning loops embedded in its cars to complement its offering (transportation) with entertainment that passengers may enjoy during a trip.

Learning and iterations. Traditional design practices, finally, have relevant limitations in terms of learning. In fact, design-build-test iterations that fuel learning are confined within a project. They are discontinued once a product is released. New learning that comes from the observation of real use can only feed the development of future versions. Innovation therefore happens episodically, in lumps. And as the context evolves new solutions became rapidly “old.”

AI drastically removes limits in learning. Note that AI factories are intrinsically iterative. They deliver through loops. As the case of Netflix illustrates, each time a customer accesses the service, the firm activates a problem-solving loop. This loop not only leverages the most recent data and algorithms. It also offers a new opportunity to further learn. The algorithm, in particular, can direct the learning strategy toward improvements, that is, toward refining its parameters to solve a problem better (e.g., showing a more appropriate movie cover to a specific user), or toward exploring new opportunities (e.g., proposing to the user a new movie category). This balancing act of exploitation and exploration, facilitated by reinforcement learning and double-armed bandit algorithms, occurs continuously, throughout the entire product life cycle.

The implications in terms of innovation are significant. First, learning *never ends*. The solution experienced by a specific user in a specific moment is not the same she experienced when the product was first released. It is the most advanced design so far. In a way, the solution is always “new.” Second, learning is *based on real use*. Rather than coming from testing prototypes in simplified contexts, here learning comes from the actual use of the product in a real context. Third, learning is *person centered*. Rather than leveraging insights from other people who used previous generation products (or tested a prototype), now data come from earlier use by the same person. Fourth, every user interaction is an opportunity *to conduct new experiments*. Learning loops are therefore designed with a different logic than traditional products. The latter included only the features that were considered useful

at the time of design. AI engines are instead *overloaded* by elements whose utility is not fully exploited at the time of release. In other words, they are explicitly designed with redundant affordances (Gibson, 1977), as we saw in the case of Tesla, where the internal pointing camera has not delivered any feature for two years.

In summary, AI factories incorporate and further empower the principles of design thinking: beyond being people centered, they are single-person centered; they facilitate creativity across segments, stakeholders, and industries, enabling abductions beyond the scope which a product was initially conceived for; finally, they are intrinsically iterative, moving learning and innovation beyond development into the product life cycle.

Design for AI

If AI empowers a more advanced practice of design, the converse can also happen: design can empower a more effective, human-centered implementation of AI. Think of the hospitality industry. Both Booking.com and Airbnb make intense use of AI, for example for personalized listing and helping hosts make decisions regarding pricing. Yet, Booking.com’s innovation path is less driven by design, but, rather, by an intense use of A/B testing. At Booking.com features are therefore pushed “from the lab outwards” rather than “from the user-inwards.” On the other side, Airbnb has design thinking in its DNA, as two of its founders, Brian Chesky and Joe Gebbia, are alumni of Rhode Island School of Design. In 2011 the company launched the Snow White project to bring human-centered design at all levels of the organization, and redesign competitive strategy (Fields Joffrion, 2018). The project was led by Rebecca Sinclair, then head of user experience research and design, and a former designer at IDEO. “At the time, like a lot of tech startups, we called the website and the app ‘the product,’” says Sinclair. But then “by practicing design thinking [...] we were looking at a journey [...], imagining our customers booking, and we saw that the moments that mattered most were offline. This offline experience—this trip to Paris or stay in a treehouse—is what they were buying from us, not a website or an app. That’s when we started to say, ‘the product is the trip’ and began shifting our perspective.” The result of this design perspective in driving innovation is evident not only by comparing

Airbnb's user interface with Booking.com, but also in the capability of Airbnb to funnel AI toward the development of new business categories, such as Airbnb Experiences.

Microsoft offers another insight into the key role of design for the implementation of AI. As Microsoft's CEO, Satya Nadella, stated, AI is the new "runtime" of its firm. Its operating model is now built around AI. This required the company to radically reorganize its IT and data assets, which had been dispersed across the company's various operations (Iansiti and Lakhani, 2020a). Interestingly, the transformation was not led by an IT manager or IT experts. Rather, the whole initiative was driven by Kurt Del Bene, an executive with product experience, as he was the former head of Microsoft's Office business unit, and a team of leaders and engineers from product functions. Nadella indeed wanted the company operating processes and AI factory to be designed as one designs products rather than IT infrastructures.

Implications for Innovation and Design Theories

Professor Simon ignores the possibility mentioned in my article that problem solving and problem finding might require opposite, or at least orthogonal cognitive strategies (and by 'cognitive' I mean not just rational, but emotional and motivational as well). (Csikszentmihalyi, 1988b, p. 184)

In 1988, before recent advancements in computer intelligence, and of the challenges that this posed to our understanding of cognition, Mihaly Csikszentmihalyi and Herbert Simon started a dispute on the true nature of creativity. Simon and Csikszentmihalyi were addressing a question that is central for innovation and design scholars: how do we think creatively? How do we have ideas and find solutions?

Simon, in exploring the potential of a computer program called "BACON" that he and his colleagues had developed at Carnegie Mellon University, was supporting a rational perspective of cognitive processes (Simon, 1988), where creativity could be interpreted as a process of problem solving (and therefore, partly embedded into computers). In a following paper, Csikszentmihalyi challenged this perspective (1988a, p. 160): "Simon wishes to prove, namely, that creativity is nothing but problem solving"; Csikszentmihalyi instead proposed "problem

finding as the hallmark of creativity." Simon (1988) reacted to Csikszentmihalyi's challenge by further reinforcing its position (p. 178): "I would claim that, just as finding laws that explain data is a problem-solving process, so *finding good problems* and finding relevant data for solving them *are problem-solving processes* of a normal kind" (our italics). The essence of the response by Csikszentmihalyi is in the opening statement of this section: problem solving and problem finding do have a different nature.

This dispute anticipated the evolution of innovation and design theories in the years to follow, with two rather independent streams unfolding: innovation as a process of problem solving, or innovation as a process of problem finding, or, in other words, as sensemaking. The first perspective (advocated by Simon) took the spotlight. Indeed, innovation scholars, especially those who investigated the process of product development, mainly looked at innovation as the result of creative problem solving (Clark and Fujimoto, 1991; Krishnan, Eppinger and Whitney, 1997; Ulrich and Eppinger, 1995). In this perspective innovation challenges can be described as a hierarchical tree (Clark, 1985), where solutions at a higher level become objectives for lower level problems (enacting Simon's view that problem finding can be seen as nested problem solving). This perspective has captured the larger share of attention also in the development of theories of design driven innovation, in which the d-school at Stanford and the related frameworks of Design Thinking are rooted (Buchanan, 1992; Brown, 2008, 2009; Kelley and Kelley, 2013; Martin, 2009). Although Design Thinking also embraces the framing of a problem (as for example in the double diamond model), it is still theoretically rooted in the theories of problem solving laid down by Simon (in which problem framing is still considered a rational activity *included* in problem solving).

This focus of theory development in the past decades was justified by the fact that problem solving was complex and therefore required the most significant chunk of effort by humans. However, the current diffusion of AI is dramatically changing this scenario. Problem solving is now increasingly embedded into the automated learning loops of AI factories. If problem solving is performed by machines, what kind of thinking is left to humans in innovation? The role of humans in AI factories (indicated in the phase "design" in Figure 2) becomes to understand

what problems should be addressed and to drive the continuous evolution of algorithms toward a meaningful direction. The core of this activity is not problem solving, but problem finding.

The consequences for the theories of innovation are substantial. In fact, as Csikszentmihalyi (1988b) clarifies in the dispute with Simon, “problem solving and problem finding might require opposite, or at least orthogonal cognitive strategies” (p. 184). This implies that the theoretical framework of problem solving, that we extensively leveraged in the past to understand innovation, will be less effective to understand human creativity in the context of AI. We need to complement those theories with new frameworks.

In his dispute, Csikszentmihalyi also suggested a possible path for these new frameworks, leveraging earlier studies he conducted on objects and products (Csikszentmihalyi and Rochberg-Halton, 1981): problem finding is an activity of meaning making, or, in other words, of sensemaking. Just to mention a simple example put forward by Csikszentmihalyi in his discussion with Simon: an algorithm that has been created to solve a problem cannot refuse to solve it; it cannot pull the plug (unless this trigger is already incorporated in its code). A human can. She can avoid to create, if it does not make sense, morally, emotionally, or by intrinsic motivation.

In the past, the perspective of innovation as an activity of sensemaking (i.e., of giving meaning to things and experiences) has only timidly found space in innovation studies. A few scholars, mainly in the field of design driven innovation, have plunged deeply in problem framing (see for example Dorst, 2015; Schön, 1982, 1995) and innovation of meaning (starting from Krippendorff’s [1989] definition that “Design is making sense of things”; see also Jahnke [2013]; Krippendorff [2006]; Norman and Verganti [2014]; Stigliani and Ravasi [2012]; Verganti [2008, 2009]; Verganti and Öberg [2013]). Our understanding of innovation as sensemaking is still very limited.

There is, however, a relevant body of theories, which has developed outside the circles of innovation scholars that we can leverage to address this new theoretical challenge. Sensemaking in organizations has indeed received significant attention in organizational psychology since the work of Weick that addresses how people give meaning to their collective experiences (Weick [1995]; Weick, Sutcliffe and

Obstfeld [2005]; for an extensive review see Maitlis and Christianson [2014]). Of particular interest for investigations of innovation and design is the focus on the construction of *new* meaning, also indicated as sense giving (Gioia and Chittipeddi, 1991) or sense breaking (Pendleton-Jullian and Seely Brown, 2016).

There is therefore an enormous (and intriguing) space ahead to be explored. We predict that the most significant future theoretical developments in innovation theories will come from a deeper understanding of problem finding and will leverage theories of sensemaking. Also, we predict that design will move closer to organization theories, and especially leadership, which is an inherent act of sensemaking (Scharmer, 2007).

Conclusion and Future Research Directions

The emergence of software, digital networks, and AI is driving widespread transformation across the economy. AI automates decision-making and learning, which is the core of innovation. The potential impact on innovation performance, as seen in the examples discussed in this article, is important. By removing the typical limitations (in scale, scope, and learning) of human-intensive design, AI can offer better performance in terms of customer centricity, creativity, and rate of innovation.

Yet, to capture this potential, managers need to fundamentally rethink the way their organization innovates. Design practice, in the age of AI, is completely different than the human-intensive innovation processes many organizations have in place today. For example, in AI-powered organizations, the role of humans is not to develop full solutions (which evolve in real time by AI), but to understand which innovation problems are meaningful, framing the innovation effort, and set up the software, data infrastructure, and problem-solving loops that will solve them.

In this article we have illustrated how pioneering organizations, such as Netflix and Airbnb, have implemented this new design practice, and how they use it to create value. Still, we are at the beginning of a transformation in innovation processes, whose extent is difficult to fully capture. Many fundamental questions are still open. For example: are AI-powered innovation practices appropriate in any context, or does their potential depends on industry or on company-specific factors, including for

example strategy or culture? Or, how can organizations transition from human-intensive to AI-centric innovation systems? Which changes of competences are required (for example, we showed that designing problem-solving loops requires new sets of skills), and which roles should lead this transition (the changes will reach across R&D, manufacturing, sales, IT and beyond)? As other pioneering managers and organizations will explore the adoption of AI in innovation, these questions will find new and more profound answers.

For scholars, the implications in terms of innovation and design theory are also substantial. New theoretical questions arise and new frameworks are needed. For example: how can we define and conceptualize innovation, in a context where change is never over and a solution is never “old”? We have seen, in fact, that problem-solving loops can keep learning and continue to deliver improved solutions to a user. How does one apply concepts such as incremental and radical innovation in a context in which the solution keeps evolving? Another example is the concept of accountability in innovation. We have seen that in AI factories solutions are created, improved, and personalized by machines, which operate through loops that scale up rapidly, with the potential of creating unintended outcomes, including the amplification of biases. Are existing theoretical frameworks that connect decisions to outcome in innovation still valid, when decisions are made by machines? And do current models of incentivizing and rewarding innovation still hold up?

One the most fascinating theoretical avenues, in our view, concerns the way scholars interpret decision-making in innovation. For one, as problem-solving is increasingly delegated to machines, humans will more deeply engage in problem finding (i.e., collectively defining which problems make sense to address). However, we still know little of how problem finding in innovation occurs. Past innovation theory has focused largely on problem-solving. A focus on problem finding would require new theoretical lenses. In this article we have suggested that future innovation and design frameworks could leverage theories of sensemaking. This would bring innovation even closer to organization theory, where sensemaking has been deeply explored, as with theories of leadership. One thing is for sure—this space promises to be one of the most fascinating journeys for innovation scholars in the years to come.

References

- Boland, R. J., and F. Collopy. 2004. *Managing as designing*. Palo Alto, CA: Stanford Business Books.
- Brown, T. 2008. Design thinking. *Harvard Business Review* 86 (6): 1–10.
- Brown, T. 2009. *Change by design: How design thinking transforms organizations and inspires innovation*. New York: Harper Business Press.
- Buchanan, R. 1992. Wicked problems in design thinking. *Design Issues* 8 (2): 5–21.
- Calabretta, G., and M. Kleinsmann. 2017. Technology-driven evolution of design practices: Envisioning the role of design in the digital era. *Journal of Marketing Management* 33 (3–4): 292–304.
- Carr, D. 2013. Giving viewers what they want. *New York Times*. Available at: <https://www.nytimes.com/2013/02/25/business/media/for-house-of-cards-using-big-data-to-guarantee-its-popularity.html>.
- Chandrashekar, A., F. Amat, J. Basilico, and T. Jebara. 2017. Artwork personalization at Netflix. Available at: <https://netflixtechblog.com/artwork-personalization-c589f074ad76>.
- Chang, R. 2017. Using machine learning to predict value of homes on Airbnb. Available at: <https://medium.com/airbnb-engineering/using-machine-learning-to-predict-value-of-homes-on-airbnb-9272d3d4739d>.
- Clark, K. 1985. The interaction of design hierarchies and market concepts on technological evolution. *Research Policy* 14: 235–51.
- Clark, K. B., and T. Fujimoto. 1991. *Product development performance*. Cambridge, MA: Harvard Business School Press.
- Csikszentmihalyi, M. 1988a. Motivation and creativity: Toward a synthesis of structural and energetic approaches to cognition. *New Ideas in Psychology* 6 (2): 159–76.
- Csikszentmihalyi, M. 1988b. Solving a problem is not finding a new one: A reply to Simon. *New Ideas in Psychology* 6 (2): 183–86.
- Csikszentmihalyi, M., and E. Rochberg-Halton. 1981. *The meaning of things: Domestic symbols and the self*. Cambridge, UK: Cambridge University Press.
- Dai, P. 2017. Helping guests make informed decisions with market insights. Available at: <https://medium.com/airbnb-engineering/helping-guests-make-informed-decisions-with-market-insights-8b09dc904353>.
- Dell’Era, C., C. Cautela, S. Magistretti, R. Verganti, and F. Zurlo. 2020. 4 kinds of design thinking: From ideating to making, engaging, and criticizing. *Creativity and Innovation Management*. <https://doi.org/10.1111/caim.12353>.
- Domingos, P. 2012. A few useful things to know about machine learning. *Communications of the ACM* 55(10): 78–87.
- Dorst, K. 2015. *Frame innovation: Creating new thinking by design*. Cambridge, MA: MIT Press.
- Fields Joffrion, E. 2018. The designer who changed Airbnb’s entire strategy. *Forbes*. July 9. Available at: <https://www.forbes.com/cdn.ampproject.org/c/s/www.forbes.com/sites/emilyjoffrion/2018/07/09/the-designer-who-changed-airbnbs-entire-strategy/amp/>.
- Galle, P. 2002. Philosophy of design: An editorial introduction. *Design Studies* 23 (3): 211–18.
- Gibson, J. J. 1977. The theory of affordances. In *Perceiving, acting, and knowing: Toward an ecological psychology*, ed. R. Shaw and J. Bransford, 67–82. Hillsdale, NJ: Lawrence Erlbaum.
- Gioia, D. A., and K. Chittipeddi. 1991. Sense Making and sense giving in strategic change initiation. *Strategic Management Journal* 12 (6): 433–48.
- Iansiti, M., and K. Lakhani. 2020a. *Competing in the age of artificial intelligence*. Cambridge, MA: Harvard Business Press.

- Iansiti, M., and K. Lakhani. 2020b. Competing in the age of AI. *Harvard Business Review* 98 (1): 59–67.
- Iansiti, M., and K. Lakhani. 2020c. From disruption to collision, the new competitive dynamics. *Sloan Management Review* 61 (3).
- Jahnke, M. 2013. *Meaning in the making: Introducing a hermeneutic perspective on the contribution of design practice to innovation*. Gothenburg: University of Gothenburg.
- Kelley, D., and T. Kelley. 2013. *Creative confidence: Unleashing the creative potential within us all*. New York: Crown Publishing.
- Krippendorff, K. 1989. On the essential contexts of artifacts or on the proposition that “design is making sense (of things)”. *Design Issues* 5 (2): 9–38.
- Krippendorff, K. 2006. *The semantic turn: A new foundation for design*. Boca Raton, FL: CRC Press.
- Krishnan, V., S. D. Eppinger, and D. E. Whitney. 1997. A model based framework to overlap product development activities. *Management Science* 43 (4): 437–51.
- Lambert, F. 2017. Tesla Model 3 is equipped with a driver-facing camera for Autopilot and Tesla Network. Available at: <https://electrek.co/2017/08/01/tesla-model-3-driver-facing-camera-autopilot-tesla-network/>.
- Lambert, F. 2018. Tesla will have a KITT-like artificial intelligence in its vehicles, says Elon Musk. Available at: <https://electrek.co/2018/07/29/tesla-kitt-like-artificial-intelligence-elon-musk/>.
- Lambert, F. 2019. Tesla reveals how it will use camera inside Model 3 to personalize in-car experience. Available at: <https://electrek.co/2019/07/24/tesla-use-camera-inside-cars-personalize-in-car-experience/>.
- Liedtka, J. 2015. Perspective: Linking design thinking with innovation outcomes through cognitive bias reduction. *Journal of Product Innovation Management* 32 (6): 925–38.
- Love, T. 2000. Philosophy of design: A metatheoretical structure for design theory. *Design Studies* 21 (3): 293–313.
- Magistretti, S., C. Dell’Era, and R. Verganti. 2020. Searching for the right application: A technology development review and research agenda. *Technological Forecasting and Social Change* 151: 119879. <https://doi.org/10.1016/j.techfore.2019.119879>.
- Maitlis, S., and M. Christianson. 2014. Sensemaking in organizations: Taking stock and moving forward. *Academy of Management Annals* 8 (1): 57–125.
- Margolin, V. 1989. *Design discourse: History, theory and criticism*. Chicago: University of Chicago Press.
- Margolin, V., and R. Buchanan. 1996. *The idea of design*. Cambridge, MA: The MIT Press.
- Martin, R. 2009. *The design of business: Why design thinking is the next competitive advantage*. Boston: Harvard Business Press.
- Mayfield, J., K. Puttaswamy, S. Jagadish, and K. Long. 2016. Data infrastructure at Airbnb. Available at: <https://medium.com/airbnb-engineering/data-infrastructure-at-airbnb-8adfb34f169c>.
- Micheli, P., S. J. Wilner, S. H. Bhatti, M. Mura, and M. B. Beverland. 2019. Doing design thinking: Conceptual review, synthesis, and research agenda. *Journal of Product Innovation Management* 36 (2): 124–48.
- Norman, D. A., and R. Verganti. 2014. Incremental and radical innovation: Design research vs. technology and meaning change. *Design Issues* 30 (1): 78–96.
- Orlikowski, W. J. 2010. Practice in research: Phenomenon, perspective and philosophy. In *Cambridge handbook of strategy as practice*, ed. D. Golsorkhi, L. Rouleau, D. Seidl, and E. Vaara, 23–33. Cambridge: Cambridge University Press.
- Pendleton-Jullian, A. M., and J. Seely Brown. 2016. *Pragmatic imagination: Single from design unbound*. San Francisco: Blurb.
- Saarinen, K. 2017. Building a visual language: Behind the scenes of our new design system. Available at: <https://airbnb.design/building-a-visual-language/>.
- Scharmer, O. 2007. *Theory U: Leading from the future as it emerges*. Cambridge, MA: The Society for Organizational Learning.
- Schleifer, A. 2017. The way we build: How rethinking the Airbnb app changed the way we approach design. Available at: <https://airbnb.design/the-way-we-build/>.
- Schön, D. A. 1982. *The reflective practitioner: How professionals think in action*. New York: Basic Books.
- Schön, D. A. 1995. *Frame reflection: Toward the resolution of policy intractable controversies*. New York: Basic Books.
- Schrage, M. 1999. *Serious play*. Boston, MA: Harvard Business School Press.
- Seidel, V. P., and S. K. Fixson. 2013. Adopting design thinking in novice multidisciplinary teams: The application and limits of design methods and reflexive practices. *Journal of Product Innovation Management* 30 (1): 19–33.
- Simon, H. A. 1982. *The sciences of the artificial*. Cambridge, MA: The MIT Press.
- Simon, H. A. 1988. Creativity and motivation: A response to Csikszentmihalyi. *New Ideas in psychology* 6 (2): 177–81.
- Spangler, T. 2018. Netflix eyeing total of about 700 original series in 2018. Available at: <https://variety.com/2018/digital/news/netflix-700-original-series-2018-1202711940>.
- Srinivasan, S. 2018. Learning market dynamics for optimal pricing: Combining elements of machine learning with structural modeling. Available at: <https://medium.com/airbnb-engineering/learning-market-dynamics-for-optimal-pricing-97effbcc53e3>.
- Stigliani, I., and D. Ravasi. 2012. Organizing thoughts and connecting brains: Material practices and the transition from individual to group-level prospective sensemaking. *Academy of Management Journal* 55 (5): 1232–59.
- Ulrich, K. T., and S. D. Eppinger. 1995. *Product design and development*. New York: Mc Graw-Hill.
- Verganti, R. 2008. Design, meanings, and radical innovation: A meta-model and a research agenda. *Journal of Product Innovation Management* 25 (5): 436–56.
- Verganti, R. 2009. *Design-driven innovation—Changing the rules of competition by radically innovating what things mean*. Boston: Harvard Business Press.
- Verganti, R. 2011a. Designing breakthrough products. *Harvard Business Review* 89 (10): 114–20.
- Verganti, R. 2011b. Radical and technology epiphanies: A new focus for research on design management. *Journal of Product Innovation Management* 28 (3): 384–88.
- Verganti, R., and A. Öberg. 2013. Interpreting and envisioning: A hermeneutic framework to look at radical innovation of meanings. *Industrial Marketing Management* 42 (1): 86–95.
- Weick, K. 1995. *Sensemaking in organisations*. London: Sage.
- Weick, K., K. M. Sutcliffe, and D. Obstfeld. 2005. Organizing and the process of sensemaking. *Organization Science* 16 (4): 409–21.
- Wilson, M. 2015. Inside the hip redesign of Holiday Inn Express: Ideo helps the budget hotel chain rethink everything from beds to mood lighting. Available at: <https://www.fastcompany.com/3043478/inside-the-hip-redesign-of-holiday-inn-express>.