

ORGANIZATIONS AS ARTIFICIAL INTELLIGENCES: THE USE OF ARTIFICIAL INTELLIGENCE ANALOGIES IN ORGANIZATION THEORY

FELIPE A. CSASZAR
University of Michigan

TOM STEINBERGER
Korea Advanced Institute of Science & Technology (KAIST)

A rarely acknowledged fact about organization theory (OT) is that many of its ideas stem from the field of artificial intelligence (AI). For example, key OT concepts such as problemistic search, heuristics, exploration, requisite variety, and organizational scripts all have their roots in AI. The main goal of this paper is to expose the full range of AI ideas that have been used in OT. We do so by explaining key AI ideas and showing how OT used them. Our review covers over 100 OT works that depend on AI ideas both critically and explicitly. We group these ideas into 10 AI approaches that speak to three fundamental processes in organizations: search, representation, and aggregation. We argue that this broad and deep borrowing from AI stems from fundamental structural similarities between AI and OT, as both fields study how artificial systems (programs and organizations) can pursue intelligent behavior. We also identify areas of AI from which OT scholars may continue to draw inspiration and suggest ways in which AI technologies may continue to affect organizations. Overall, our work shows that, beyond its effect as a technology, AI has given OT a set of models about how organizations work.

It is well known that organization theory (OT) has borrowed many ideas from other disciplines. For example, organizational ecology, evolutionary economics, and fitness landscapes have roots in ideas from evolutionary biology (Hannan & Freeman, 1989; Levinthal, 1997; Nelson & Winter, 1982). Other examples include ideas borrowed from sociology, such as social construction (Zajac & Westphal, 2004), institutional logics (DiMaggio & Powell, 1983), and structural holes (Burt, 1992), and ideas borrowed from psychology, such as categorization (Pontikes, 2018), judgment biases (Lovallo & Kahneman, 2003), and personality traits (Hambrick & Mason, 1984).

Perhaps less known is OT's fruitful borrowing from the field of artificial intelligence (AI). OT concepts such as search, heuristics, and representation, to name a few, have roots in AI. Such borrowing is not coincidental. For one thing, AI and OT both study complex systems and, hence, concepts such as

adaptation, hierarchy, and information apply to both (Axelrod & Cohen, 2001: 1–31). For another, AI and OT are both “sciences of the artificial” (Simon, 1969/1996), investigating systems that adapt to their environment to fulfill goals, so that concepts such as coordination, feedback, and learning apply to both. More generally, borrowing from AI to understand organizations is natural because the technologies of the day have always served as metaphors for the most complex systems (Morgan, 2006: 12). For instance, the brain has been compared in its time to the hydraulic pump (Descartes, 1633/1985), the steam engine (Freud, 1933), and the computer (von Neumann, 1958).

Acknowledging the AI roots of many OT ideas has at least two important benefits for OT. The first is practical: being aware of the range of AI ideas used by OT can help us better understand how AI itself will affect organizations, a topic of much current interest (Bailey, Faraj, Hinds, von Krogh, & Leonardi, 2019: 642; Baum & Haveman, 2020: 270–271). For an example of the cost of lacking that understanding—that is, the cost of a narrow framing of how AI affects organizations—consider the popular view that the main effect of AI on organizations is to decrease the cost of making predictions (Agrawal, Gans, &

The authors are grateful to Linda Argote and two anonymous referees for their insightful comments and guidance. For many helpful comments, the authors thank Dan Levinthal, Cha Li, and Jim Walsh. All errors remain the authors' own.

Goldfarb, 2018). While this effect obviously exists, understanding AI strictly through the lens of prediction costs can miss other important effects of AI on organizations, such as its effect on the type of knowledge the organization can represent and on the organization's search and aggregation processes.

The second benefit is theoretical: knowing the AI roots of current organizational theories allows us to create better organizational theories. Organizations can be viewed as artificial systems whose main purpose is to pursue intelligent behavior (March, 1999; Ocasio, Rhee, & Boynton, 2020). In other words, organizations can themselves be viewed as a type of artificial intelligence (Csaszar, 2018: 608). Being able to borrow from a field entirely devoted to the development of intelligent systems is therefore an enormous benefit, akin to medicine's fruitful (to put it mildly) borrowing from chemistry once it was understood that the human body is, among other things, a chemical system. Today, for example, many in OT understand AI only or mainly from its most popular current manifestation: machine learning. Such an understanding, however, misses the many productive analogies that have been established between AI and OT over the past 60 years. Such a circumscribed view is not just unscholarly but also risky, as ignoring these links limits the repertoire of ideas available to organizational theorists.

The overarching goal of this paper is to help the OT field obtain these two benefits, one practical and one theoretical. The bulk of our paper is therefore devoted to uncovering the connections between AI and OT. We do this by showing how seminal papers in OT explicitly borrowed specific ideas from AI (in a few places, we also document the reverse influence: cases in which AI has borrowed from OT). We group the AI-OT linkages into 10 "approaches" to AI—fundamentally different ways of achieving AI. For each of the 10 areas of connection between AI and OT that we uncover, we describe the seminal AI idea and elaborate on its use by multiple authors in OT. For the sake of clarity, we group the 10 AI approaches into three themes: search, representation, and aggregation. Note that, because both the AI literature and OT's use of it are vast, our paper is necessarily incomplete. Our review of these literatures should be seen as illustrative of the connections, but not exhaustive.

This paper makes three main contributions to the literature. First, we expose and clarify the broad range of AI ideas that have been put to use in OT. Here, the multiplicity of borrowings is as important as the specific ideas borrowed. Second, we identify areas of AI from which OT scholars may continue to draw

inspiration. We do this by identifying AI ideas overlooked by OT that may illuminate classic OT questions and by showing how AI ideas that were imported into OT have diverged since the import took place and may therefore be ripe for new borrowing. Third, we suggest ways in which AI technologies may continue to affect organizations. We do this by devising plausible scenarios about the future use of AI in organizations and by building on recent AI ideas that aim to limit the potentially negative effects of AI technologies.

The remainder of the paper is structured as follows: the next section provides the background necessary to understand the bridges that have been established between AI and OT. The section after that presents the methodology used to review the literature. The three ensuing sections—focused, respectively, on search, representation, and aggregation—expose the multiple ways in which OT has borrowed from AI. The final section discusses how AI may continue to affect the theory and practice of organizing.

BACKGROUND

To contextualize how OT has borrowed from AI, in this section, we define AI and its origins, describe the three themes that serve as an organizing framework for our work, and show that some AI ideas have indeed originated in OT.

Defining AI

There is, in fact, no generally accepted definition of AI. Russell and Norvig (2020: 1–4) aptly arranged the multiple existing definitions along two axes—(a) "thinking" versus "acting" and (b) "humanly" versus "rationally"—which creates the four categories of definition shown in Figure 1:

1. AI as devising algorithms that *think humanly*. This is also known as "cognitive modeling" and its goal is to better understand human cognition.

FIGURE 1
Four Quadrants of AI Definitions

	Humanly	Rationally
Thinking	Think Humanly (e.g., cognitive modeling)	Think Rationally (e.g., logic)
Acting	Act Humanly (e.g., Turing test)	Act Rationally (e.g., engineering)

2. AI as devising programs that *think rationally*. This requires using rational processes such as logic and probability.
3. AI as devising machines that *act humanly*—that is, machines that could pass the Turing test by credibly impersonating a human being.
4. AI as devising programs that *act rationally*. This is consistent with an engineering approach: what matters is acting optimally with respect to a given measure of performance, regardless of the specific process used. That is, there is no need to think like humans or to use well-sanctioned logic; the goal is just to perform in a way that would normally require intelligence.

Following Russell and Norvig (2020), we include within the scope of AI all research that fits within any of these four quadrants. We believe it is useful for OT to use a broad definition of AI, as this avoids missing AI–OT connections. To illustrate this point, it is useful to compare OT to cognitive science in terms of their borrowing from AI. For cognitive science, which studies the computational processes used by the human brain, it makes sense to limit its borrowing from AI to the “think humanly” quadrant. However, because OT studies organizations (which typically employ multiple humans and technologies), restricting the borrowing to a specific quadrant would be counterproductive, as not all information processing in organizations is limited by the constraints of one human brain. In other words, because AI contains the models from cognitive science and because not all components of organizations are human, AI is a richer source of OT analogies than cognitive science.

The Origins of AI

The idea of thinking beings created by humans—and the practical and philosophical problems that might arise if they existed—has a long prehistory. There are, for example, the Greek myths of bronze robots made by the smith-god Hephaestus (Mayor, 2018) and the Jewish folklore of clay-based robots called “golems” (McCorduck, 2004: 14–15). Also contributing to the eventual development of AI was the development of the mathematics of functions commonly associated with thinking—particularly, logic and probability. Key steps in the development of logic were Aristotle’s (c.350 BCE/1984) work on syllogisms, Boole’s (1854/1951) on propositional logic, and Frege’s (1879/1967) on predicate logic. Key steps in the development of probabilistic

thinking were Pascal’s and Fermat’s work on computing probabilities, Bernoulli’s work on utility, and Bayes’s work on how to rationally use new data to update probabilities (Hacking, 2006). Finally, there were the first attempts to create an actual mechanical computer, made by Babbage and Lovelace starting in the 1830s. This work made previous speculations about AI more tangible and raised questions that are still important. For example, in what accounts for an instruction manual for their planned computer, Lovelace (1843) discussed whether machines would ever be intelligent.

For another century, the field of AI continued to consist almost exclusively of philosophical investigation. The current history of AI begins with the creation of the first electronic computers in the 1940s. This technology made it possible to try out the various ideas about how to create intelligent machines and see how far they could take us. The result was something of an arms race—or a Cambrian explosion—of methods to produce AI. The most important of these methods will be presented when we review the linkages between AI and OT.

The term “artificial intelligence” was first used by McCarthy, Minsky, Rochester, and Shannon (1955/2006) in their proposal for what became the Dartmouth Summer Research Project on Artificial Intelligence, a summer-long workshop during which the main AI researchers at the time could share their ideas. The stated goal of the conference served to define the field: “to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it” (McCarthy et al., 1955/2006). Some 20 researchers attended at least part of this seminal event, including Herbert Simon, Allen Newell, Claude Shannon, Marvin Minsky, John McCarthy, John Holland, Ross Ashby, Warren McCulloch, John Nash, and Arthur Samuel.

Three Themes in AI Research: Search, Representation, and Aggregation

The work of these and other pioneers produced many ideas on how to achieve AI and many groupings of these ideas have been proposed, as can be seen in the tables of content of the main AI textbooks (Luger & Stubblefield, 1993; Nilsson, 1998; Rich, 1983; Russell & Norvig, 2020; Winston, 1977). We chose our particular grouping into 10 approaches because it is useful for our purpose of elucidating AI–OT connections. For the sake of clarity and to

better appreciate how these AI approaches have been used within OT, we group these 10 approaches into three themes: (1) search, (2) representation, and (3) aggregation.

“Search” approaches conceptualize AI in terms of a problem and some means to solve it. For a Roomba—a commercial robot that vacuums the floor—the goal is to clean a room and the means are the Roomba’s possible actions, such as moving forward, turning to the right, and backing up. Different search approaches use different techniques to generate alternatives and to pick among those.

“Representation” approaches conceptualize AI in terms of a given task environment and some means to model it (so that an AI can act on it). For example, developing AI to manage customer comments would require devising a model to distinguish, among other things, between positive and negative emotions in written comments. Different representation approaches would have different ways of encoding such models.

“Aggregation” approaches conceptualize AI in terms of strategies for combining several subsystems to create a desired aggregate behavior. For example, if different navigation systems in an autonomous car—camera, radar, radio contact with other cars—recommend different actions—say, stopping versus slowing down versus swerving—what would be the best way to combine these recommendations? Different aggregation approaches would propose different combination strategies.

The Contribution of OT to AI

Although this paper emphasizes the influence of AI on OT, it’s worthwhile to note that there are many cases of influence in the other direction. For example, Herbert Simon attributes his AI work on *heuristic problem-solving* (Approach 2, according to the scheme we will present later) to previous work that he had done on organizations to better understand how humans and organizations make decisions (Spender, 2013: 329). Edward Feigenbaum’s work on *expert systems* (Approach 5) was influenced by his early experience as a research assistant for Herbert Simon and Richard Cyert while they were modeling organizational decision-making processes and had to extract business knowledge from managers (Feigenbaum, 1992: 195). The literature on *distributed AI* (Approach 9) builds on classic models of coordination, power, and trust in organizations (Weiss, 2013: xxxvi).

METHODOLOGY

To review how OT has built on AI ideas, we used the following methodology:

1. Create a list of AI keywords. We consolidated the topics covered by the principal AI textbooks from the last 45 years (Charniak & McDermott, 1985; Luger & Stubblefield, 1993; Nilsson, 1998; Poole & Mackworth, 2010; Rich, 1983; Russell & Norvig, 2020; Winston, 1977). We included not just recent AI textbooks but also older ones, as these also may include ideas cited in OT. Keywords included, for example, “artificial intelligence,” “machine learning,” “neural network,” “expert system,” and “representation.”
2. Search the main bibliographic databases (JSTOR, ISI Web of Science, and Google Scholar) for use of those keywords in the main OT and strategy journals (*Academy of Management Journal*, *Academy of Management Review*, *Administrative Science Quarterly*, *Journal of Management*, *Journal of Management Studies*, *Management Science*, *Organization Science*, *Organization Studies*, *Strategic Management Journal*, and *Strategy Science*).
3. Eliminate from the list those articles that contain only superficial citations of AI ideas.
4. Read the remaining articles and code them in terms of the main AI ideas used.
5. Expand the list by including the forward and backward citations of the most-cited articles.
6. Repeat steps 2–5 until we are confident that we have identified relevant and representative articles illustrating the main uses of AI ideas in OT.
7. Check with our personal networks that we are not missing important works.

Three previously published review papers that are related to ours are Joseph and Gaba (2020), on organization structure and information processing; Puranam, Stieglitz, Osman, and Pillutla (2015), on modeling bounded rationality; and Baumann, Schmidt, and Stieglitz (2019), on search in rugged landscapes. Like ours, these papers are broadly interested in information-processing explanations of organizational behavior. Unlike Puranam et al. (2015) and Baumann et al. (2019), however, our focus is not on surveying modeling papers, but on OT theories (regardless of research method) that have borrowed ideas from AI. Unlike Joseph and Gaba (2020), which looked at papers related to organization structure and information processing, our focus is on OT papers

that use AI ideas. None of the above papers delves into the AI roots of the OT theories that we cover.

AI IDEAS USED IN OT: SEARCH APPROACHES

Our description of the ways in which OT has borrowed from AI follows the structure presented in Table 1. Each row in that table corresponds to an AI approach, with the left and right columns containing, respectively, AI ideas and illustrative uses in OT. Table 1 is thus a type of Rosetta Stone, connecting AI ideas to their OT counterparts.

In this section, we describe the first four rows of that table; that is, the four approaches that relate to *search*. In each subsection, we describe an AI approach—the definition, examples, and historical milestones for context—and then the ways in which that approach has been used in OT.

Approach 1: Cybernetics and Control Theory

AI idea. “Cybernetics” is the study of systems that use feedback to automatically control their behavior in order to achieve a goal (Wiener, 1948). Today, this area of research is more commonly known as “control theory.” The challenge that a cybernetic system is trying to address is to maintain an optimal behavior in a dynamic environment by changing a set of parameters, the canonical example being a heating system that uses a thermostat to keep the temperature of a room at a desired level. In the case of the thermostat, the changing parameter could be whether to keep the boiler on or off.

The impetus to study cybernetics arose between the First and Second World Wars with the need to develop anti-aircraft defenses (Mindell, 2002), which required reacting to feedback more quickly than any human could do. The rigorous study of cybernetics begins with the work of Norbert Wiener at MIT, who coined the term and developed the mathematics of cybernetics, based on dynamic systems—that is, systems of differential equations. Building on Wiener’s work, Forrester expanded the use of dynamic systems to model the behavior of complex systems such as the economy and the climate, producing ideas that inspired the creation of the Club of Rome (Meadows & Club of Rome, 1972). This line of work continues today in the System Dynamics Group at MIT (see, e.g., Sterman, 2000).

Shannon’s (1948) communication theory, which underlies all communication systems, owed much of its formulation to Wiener’s (1948) conceptualization of communication (Shannon & Weaver, 1949/

1998: 85); hence, communication theory is sometimes considered a descendant of cybernetics. Between 1946 and 1953, the Macy Foundation sponsored annual conferences on cybernetics, which were instrumental in creating the field and in diffusing its ideas (Pias, 2016).

The general applicability of the cybernetic concepts of information, communication, and feedback led other researchers to apply cybernetic ideas to several fields, including cognition (Ashby, 1952, 1956), anthropology (Bateson, 1972), biology (Maturana & Varela, 1980), and management (Ansoff, 1965; Beer, 1972).

OT uses. The influence of cybernetics in OT is pervasive. One area in which the imprint is clear is the research on aspiration levels, which has been an influential and productive area of OT research (Greve, 1998; Posen, Keil, Kim, & Meissner, 2018). March and Simon (1958/1993: 68) introduced the idea that a firm’s behavior depends on whether its performance is above or below an aspiration level, which directly parallels cybernetics’ simplest model: the thermostat. In fact, Simon was deeply aware of Wiener’s work (Dasgupta, 2003: 697) and March and Simon (1958/1993: 65) used the same type of mathematics and block diagrams found in Wiener (1948). Simon’s (1952) grasp of cybernetics is clearest in his paper “On the Application of Servomechanism Theory in the Study of Production Control.”

Ideas from cybernetics also entered into OT through the influence of communication theory. Cyert and March (1955: 130–131), for example, built on Shannon (1948) to conceptualize organizational structure as a “communication pattern” and the role of the organization designer as someone in charge of the “design of informational channels,” who must pay special attention to how the organization “receives, decodes, encodes, and retransmits information.” All of this is reminiscent of the types of problems Shannon’s theory addressed. Another early paper connecting OT to cybernetics is Bavelas’s (1950) model of the communicational efficiency of different group structures, initially presented at the 1950 Macy conference (Pias, 2016). The imprint of communication theory is also visible in the ensuing works of the Carnegie tradition (e.g., Cyert & March, 1963; Galbraith, 1973; Simon & Newell, 1958; Thompson, 1967), which view organizations as information-processing systems—a view that continues to be a very active area of research today (see, e.g., the recent survey by Joseph & Gaba, 2020).

Another cybernetic idea that has become important in OT is the idea of “requisite variety” (Ashby,

TABLE 1
10 Approaches to AI and Examples of the OT Ideas They Inspired

AI Idea	Uses in OT
Search	
1. <i>Cybernetics and control theory</i> (Ashby, 1956; Shannon, 1948; Wiener, 1948)	<ul style="list-style-type: none"> • Aspirations (Greve, 1998; March & Simon, 1958/1993; Posen et al., 2018) • Communication processes (Bavelas, 1950; Galbraith, 1973; Joseph & Gaba, 2020; Thompson, 1967) • Requisite variety and simple rules (Bingham & Eisenhardt, 2011; Siggelkow, 2002; Weick, 1979) • System dynamics and vacillation (Forrester, 1961; Gary, Kunc, Morecroft, & Rockart, 2008; Nickerson & Zenger, 2002; Sterman, 2000)
2. <i>Heuristic problem-solving</i> (Newell et al., 1959; Newell & Simon, 1956; Shannon, 1950)	<ul style="list-style-type: none"> • Boundedly rational search (Cyert, Feigenbaum, & March, 1959; Newell & Simon, 1972; Simon, 1955) • Rugged landscapes (Baumann et al., 2019; Levinthal, 1997) • Organizational programs and routines (Cohen & Bacdayan, 1994; Feldman & Pentland, 2003; Nelson & Winter, 1982)
3. <i>Evolutionary computation</i> (Holland, 1975; Koza, 1992; Turing, 1950)	<ul style="list-style-type: none"> • Exploration and exploitation (Fang, Lee, & Schilling, 2010; Laureiro-Martinez Brusoni, Canessa, & Zollo, 2015; March, 1991; Posen & Levinthal, 2012) • Organizational evolution (Bruderer & Singh, 1996; Fang et al., 2010; Lee, Lee, & Rho, 2002) • Modularity (Baldwin, 2018; Baldwin & Clark, 2000; Rivkin, 2000; Siggelkow & Levinthal, 2003)
4. <i>Reinforcement learning</i> (Bellman, 1957; Samuel, 1959; Sutton & Barto, 2018)	<ul style="list-style-type: none"> • Ambiguity in organizational learning (Denrell & March, 2001; Lave & March, 1975) • Credit assignment and model-based organizational learning (Denrell et al., 2004; Fang & Levinthal, 2009; Rahmandad, 2008) • Complex and interactive learning processes (Puranam & Swamy, 2016; Rahmandad, 2008)
Representation	
5. <i>Expert systems and knowledge representation</i> (Buchanan, Sutherland, & Feigenbaum, 1969; Dreyfus, 1972; McCarthy & Hayes, 1969/1981; Schank & Abelson, 1977)	<ul style="list-style-type: none"> • Codifying organizational knowledge (Burton & Obel, 2004; Hannan et al., 2007) • Limits of codified knowledge (Brown & Duguid, 1991; Kogut & Zander, 1992; Orr, 1996; Prietula & Simon, 1989; Suchman, 1987) • Scripts, skills, and routines (Gioia & Poole, 1984; Nelson & Winter, 1982; Pentland, 1995)
6. <i>Connectionism and machine learning</i> (Geman, Bienenstock, & Doursat, 1992; McCulloch & Pitts, 1943; Rumelhart, Hinton, & Williams, 1986)	<ul style="list-style-type: none"> • Interactive learning in games (Marchiori & Warglien, 2008) • Interpretive processes (Gavetti & Warglien, 2015) • Representational complexity (Csaszar & Ostler, 2020)
7. <i>Bayesian networks</i> (Pearl, 1988, 2000; Spirtes, Glymour, & Scheines, 2000; Wright, 1921)	<ul style="list-style-type: none"> • Strategic decision-making process (Durand & Vaara, 2009) • Dealing with causal ambiguity (Ryall, 2009) • Causal understanding (Bettis & Blettner, 2020)
Aggregation	
8. <i>Cellular automata and emergence</i> (Gardner, 1970; Varela, Maturana, & Uribe, 1974; von Neumann & Burks, 1966; Wolfram, 2004)	<ul style="list-style-type: none"> • Segregation (Schelling, 1969, 1971) • Evolution of cooperation (Nowak & May, 1992) • Population ecology dynamics (Lomi & Larsen, 1996) • Social autopoietic systems (Luhmann, 1990)
9. <i>Distributed AI</i> (Minsky, 1986; Selfridge, 1959; Star, 1989; Winograd & Flores, 1986)	<ul style="list-style-type: none"> • Organizational cognition (Carley & Gasser, 1999; Prietula et al., 1998) • Coordination and organizational processes (Malone & Crowston, 1991; Malone, Crowston, & Herman, 2003) • Transactive memory (Ren & Argote, 2011; Wegner et al., 1985) • Boundary objects (Carlile, 2002)
10. <i>Ensemble methods</i> (Breiman, 1996; Ho, 1995; Moore & Shannon, 1956/1993; Schapire, 1990)	<ul style="list-style-type: none"> • Information aggregation in organizations (Csaszar, 2013; Csaszar & Eggers, 2013) • Reliability of decision-making structures (Christensen & Knudsen, 2010) • Wisdom of the crowd (Csaszar, 2019; Grushka-Cockayne, Jose, & Lichtendahl, 2017) • Cognitive diversity (Arthur, 1994; Page, 2007, 2018)

1956); namely, that, for a system to deal successfully with an environment, it needs to have as many degrees of freedom as the environment. This idea is central in works such as Weick (1979), Burton and Forsyth (1986), Siggelkow (2002), and Csaszar and Ostler (2020).

The mathematical techniques developed by Wiener and Forrester—today, packaged in user-friendly software like iThink (www.iseesystems.com)—have been productively used in the context of organizations by the system dynamics literature. One of the main goals of this literature has been to elucidate the complex nonlinear behavior that emerges when feedback loops and accumulation processes are combined, as they often are in organizations (Gary et al., 2008). Sterman (2000) delivered a textbook treatment of this literature; examples of research papers in this tradition include Repenning (2002), Rahmandad (2008), and Freeman, Larsen, and Lomi (2012).

Another theory about organizations with a clear cybernetic origin is vacillation theory (Nickerson & Zenger, 2002), which proposes that organizations can approximate an optimal configuration by alternating between discrete states on the basis of feedback. Like aspiration models (March & Simon, 1958/1993), Nickerson and Zenger's (2002) model was based on the mathematical model of a thermostat, a connection they acknowledged multiple times in their paper.

Approach 2: Heuristic Problem-Solving

AI idea. Heuristic problem-solving is an AI approach based on using “heuristics”—rules or strategies developed through experience—to search for a sequence of actions leading from the current state to a desired state (Winston, 1992: 53). Because a great variety of problems can be represented as a sequence of states and transitions between these states, this approach can be used to solve many problems that previously were only solvable by humans.

A classic example of such a problem is the game of chess, in which one searches for a sequence of actions (moves of the chess pieces) that will lead from the current state of the board to the goal state—a checkmate. In brief, chess programs work by assigning a score to each considered state and then picking a next move that gets the player closer to a high-score state. The fact that every other move is played by the opponent is dealt with by using a “minimax” algorithm, which assumes that, at each step, the opponent's moves will try to minimize the player's expected score (which the player is trying to

maximize). The complexity of computer chess arises because the number of possible states is usually too large even for a computer; it has been estimated that, starting from the initial board, the number of reachable positions is 10^{46} (Chinchalkar, 1996). No matter how fast the computer is, a chess program can only explore a tiny fraction of the search space. Thus, except in the final stages of a game, a chess program typically cannot know which path will lead to winning the game, so it must resort to heuristics to assign a score to any given board state and to select which move to make next. Heuristics in this context could be having more valuable pieces than the opponent and having the king in a safe position.

The idea of representing problems as states and transitions between states was introduced in von Neumann and Morgenstern's (1944) seminal book on game theory. The state-and-transitions representation corresponds to their extensive-form game representation, which is usually drawn as a “tree” in which each node is a state and each link a transition. In AI, this representation is sometimes called the “game tree,” “search tree,” or the “search space.” The idea of using heuristics to prune the search space was influenced by Polya's (1945) nontechnical description of heuristics that could be used to solve mathematical problems, which Newell (1983) knew from having taken an undergraduate course with him around the time Polya's book was published.

Significant demonstrations of the power of heuristic problem-solving came with the Logic Theorist program (Newell & Simon, 1956), which was able to automatically prove several theorems from Whitehead and Russell's (1910) *Principia Mathematica*, and the General Problem Solver (Newell, Shaw, & Simon, 1959), which expanded the reach of the heuristic approach to other logical puzzles.

Early chess-playing programs were developed by Shannon (1950) and Newell, Shaw, and Simon (1958). The successes of these early examples of the heuristics approach led to overly optimistic predictions. Herbert Simon and Allen Newell (1958: 7) predicted that, “within 10 years, a digital computer will be the world's chess champion.” Simon (1965: 96) later predicted that “machines will be capable, within 20 years, of doing any work a man can do.” Marvin Minsky (1967: 2) predicted that, “within a generation ... the problems of creating ‘artificial intelligence’ will be substantially solved.”

The 1960s did not bring such triumphs, but did bring the discovery of important ideas about how to efficiently prune search spaces, including alpha-beta pruning, which cuts branches of the search tree that

are assured not to contain good solutions (Edwards & Hart, 1963), and the A* algorithm, which can be used in problems wherein there is a proxy for the value of each state (Hart, Nilsson, & Raphael, 1968). In 1997, a computer finally did defeat a world chess champion, using alpha–beta pruning and taking advantage of the exponential increase in computing power (Campbell, Hoane, & Hsu, 2002). Simon and Newell, though off by about 30 years, had been correct in their prediction that digital chess players would surpass humans.

OT uses. Several OT ideas revolving around the concept of bounded rationality were influenced by the AI work on heuristic problem-solving, a fact that is attributable to Herbert Simon having been a central figure in both OT and AI. The OT idea that individuals are “boundedly rational”—that, however rational their thinking is, they can only take so much into account—is akin to the idea that, for any sufficiently complex problem, even a computer cannot explore all the states in a search space. Simon wrote the foundational papers on bounded rationality (Simon, 1955, 1956) while he was working with Newell on what became the Logic Theorist (Newell & Simon, 1956) and the General Problem Solver (Newell et al., 1959) programs.

Because of bounded rationality, an individual cannot seek an optimal solution, but, rather, can only search for a “good enough” or “satisficing” solution—much in the same way that a computer chess program would settle for the best position found after searching a limited subset of the impossibly large search tree. To establish how such a process of search would look like in real organizations, Cyert et al. (1959) developed a behavioral simulation of the pricing process in a duopoly. Cyert and March (1963: chap. 6) later described with painstaking accuracy the pricing process of a department store.

Because, according to this view, individuals make decisions by searching, it was essential to understand how they search; that is, what search heuristics individuals use. This line of research launched a vast literature using protocol analysis (Ericsson & Simon, 1992) to examine how novices and experts solved problems. This body of knowledge is summarized in Newell and Simon’s (1972) magnum opus, *Human Problem Solving*. A descendant of this literature in entrepreneurship is the concept of “effectuation” (Sarasvathy, 2008), which characterizes how entrepreneurs discover opportunities. Sarasvathy, herself a former PhD student of Simon’s, has defined effectuation in terms of search (Sarasvathy, 2008: 65–95) and investigated it using think-aloud protocols.

Cyert and March (1963) argued that, once a firm finds a solution—a path from an initial state (say, the introduction of a new product) to a desired state (profitable sales of that product)—this set of actions becomes a standard operating procedure. From then on, these procedures can evolve in different ways, a process that has been explored by the literature on routines (Cohen & Bacdayan, 1994; Feldman & Pentland, 2003; Nelson & Winter, 1982). In current research, the process of search is a central behavioral foundation of the theoretical and empirical literature on rugged landscapes (e.g., Levinthal, 1997; see Baumann et al., 2019, for a survey of over 70 studies in this literature).¹

Approach 3: Evolutionary Computation

AI idea. “Evolutionary computation” is an AI approach based on searching a population of candidate solutions by simulating the evolutionary processes of variation, selection, and retention (Mitchell, 1996: 2). Although this approach is sometimes known as “genetic algorithms,” we will reserve that label for a specific type of evolutionary computation.

The goal of evolutionary computation is to discover high-quality solutions in vastly large and complex search spaces. In contrast to heuristic problem-solving, which searches a space starting from the current configuration, evolutionary computation finds satisficing solutions by creating candidate solutions through a process that simulates biological evolution. The premise is that, if biological evolution was able to “discover” millions of successful solutions to a complex problem—staying alive as an individual and as a species in a variety of changing and dangerous environments—a computer simulation of evolution should be able to discover successful solutions to complex problems. Evolutionary computation is, in fact, commonly used to find solutions to complex engineering design problems. Koza (2010), for example, documents 76 cases in which evolutionary computation has matched or improved state-of-the-art solutions devised by engineers or scientists, such as designs for electrical circuits and algorithms to optimize stock portfolios.

The first milestone in the development of evolutionary computation was Turing’s (1950) proposal for a “learning machine” that could imitate the principles of natural evolution. During the 1950s and

¹ Although the landscape analogy is rooted in evolutionary biology (Wright, 1932), the idea of how individuals and firms may search on that landscape comes from AI.

1960s, there were some promising but limited successes in trying to implement Turing's vision (Fogel, 1998). It was Holland's (1975) book that put evolutionary computation on a firm theoretical footing, and marked the beginning of the modern study of evolutionary computation. Although his main aim was to understand how biological evolution works, an important effect of this book was to popularize evolutionary computation as an alternative to heuristic methods in AI.

The specific type of evolutionary computation proposed by Holland (1975) is called the "genetic algorithm." It models mutation and crossover in a population of candidate solutions. Each solution is modeled as a binary vector of a fixed size that represents a position in a multidimensional landscape. Mutations randomly change bits of solutions, while crossover creates "offspring" that combine the information contained in the "parents." These two genetic operators are applied over and over to high-fitness candidate solutions to create successive generations of new solutions. As with its natural counterpart, when this blind evolutionary process is applied over and over, ever-more fit solutions tend to emerge.

To understand why genetic algorithms were effective, Holland (1975) theorized that changing the relative frequency of crossover and mutation made it possible to balance exploration with exploitation—that is, to balance a strategy of visiting new positions in the landscape with a strategy of remaining close to the current position. To analyze this issue, Holland (1975: § 5.1) used the mathematics of multiarmed bandits; that is, the idealized situation of a gambler choosing among different slot machines as he learns about their differing distributions of rewards. Koza (1992), a student of Holland at the University of Michigan, extended the genetic algorithm to allow for solutions of variable length; rather than use fixed size vectors, he used Lisp expressions. This type of evolutionary computation was called "genetic programming," as it makes it possible to "evolve" computer programs.

OT uses. Ideas from evolutionary computation entered forcefully into OT through the work of James March. Building on Holland (1975), March (1991) proposed that organizations, too, faced an exploration/exploitation trade-off. This idea became a staple of research on the structural determinants of innovation (e.g., Raisch & Birkinshaw, 2008) and of research on innovation more generally. March's borrowing from evolutionary computation is not just conceptual, as the model he develops in that paper

is a type of genetic algorithm over a population of individuals, each described by a vector of beliefs. This population of individuals is subject to variation forces (due to hiring and learning) and selection forces (based on a fitness function that depends on the accuracy of the individuals' beliefs).² Genetic algorithms also figure prominently in models of organizational evolution, such as those of Bruderer and Singh (1996) and Lee et al. (2002), and in models that extend March (1991), such as that of Fang et al. (2010).

Evolutionary computation also affected OT through the use of the multiarmed bandit to examine exploration and exploitation. Posen and Levinthal (2012), for example, analyzed organizational exploration and exploitation by using a bandit model; they studied how the optimal degree of exploration depends on factors such as environmental turbulence and the decision-maker's knowledge. Bandit models have also spawned empirical research in OT. Laureiro-Martínez et al.'s (2015) laboratory study, for example, investigated how individuals behave when making decisions involving risk.

Evolutionary computation has also affected OT through its effect on the influential literature, initiated by Baldwin and Clark (2000), on modularity and design rules. Baldwin and Clark (2000: 10) acknowledged that the organizing framework of their book is the evolutionary understanding of complex systems stemming from Holland's (1975) work. Among the AI concepts they borrowed are the idea of seeing design as search, the interpretation of evolutionary operators such as mutation and crossover, the criteria for selecting powerful evolutionary operators, and the "credit assignment problem"—that is, identifying the contribution to fitness of specific elements of a design (Baldwin & Clark, 2000: 129–130, 225, 273). Much work has built on Baldwin and Clark (2000), including research on knowledge recombination (Brusoni, Prencipe, & Pavitt, 2001), product modularity (Ethiraj & Levinthal, 2004), open innovation (Schilling, 2000), imitation (Rivkin, 2000), organizational structure (Siggelkow & Levinthal, 2003), and ecosystems (Baldwin, 2018).

² A paper that could be considered an antecedent of March (1991) is Cohen's (1981) model of parallel thinking in organizations, which borrowed the idea of modeling individuals as vectors of beliefs from Holland (1975). Cohen, both a colleague of Holland at the University of Michigan and a coauthor of March's, played an instrumental role in creating this AI-OT bridge (D. Levinthal, personal communication, December 30, 2020).

Approach 4: Reinforcement Learning

AI idea. “Reinforcement learning” is an AI approach based on learning what actions are most appropriate through an interactive process of trial and error (Sutton & Barto, 2018: 1–2). This approach captures the challenge of learning to play a game you do not know simply by playing it over and over and being told each time who won (Russell & Norvig, 2020: 789). A key challenge in such a problem is the issue of credit assignment: which of the many actions in the game contributed to the victory? Reinforcement learning is particularly helpful in dealing with problems in which heuristics are hard to acquire or in which it is unclear how to assign fitness values to different configurations. An example of reinforcement learning is the software that learned to play classic videogames only by looking at the screen while controlling the movement of the joystick (Mnih et al., 2013).

A number of important ideas in reinforcement learning come from classic models of learning from psychology, such as the law of effect (Thorndike, 1911), conditioning (Pavlov, 1927; Skinner, 1938), and Hebbian learning (Hebb, 1949). A stylized version of these ideas was first put into mathematical form by Bush and Mosteller (1955). Another important source of ideas in reinforcement learning is “dynamic programming” (Bellman, 1957), a mathematical framework to solve multistage optimization problems using backward induction. Bellman (1957: ix) coined the term “curse of dimensionality” to denote that the computational resources needed to solve learning problems increase exponentially with the size of the problem.

The use of reinforcement learning in AI starts with Samuel’s (1959) checkers-playing program, which changed the coefficients of a scoring polynomial depending on feedback it gained from playing against a copy of itself (a technique also known as “self-play”). This polynomial, representing the computer’s “understanding” of the game, is used to evaluate the quality of different board configurations. Samuel’s program was ahead of its time in that it combined multiple approaches; specifically, reinforcement learning and heuristic search (Approach 2). During the 1980s, researchers improved on Samuel’s ideas to create more effective algorithms to learn to play multistage games (examples of these algorithms, collectively called “temporal difference learning,” include TD-Lambda and Q-learning; see Sutton & Barto, 2018: 13–22 for details).

Another milestone in the history of reinforcement learning was the development of TD-Gammon

(Tesauro, 1995), the first program to play backgammon at a level similar to that of the best players of the time.³ More recently, a similar approach was used to develop AIs that beat the human champions in classic video games (Mnih et al., 2013), the ancient Chinese strategy game of Go (Silver et al., 2016), and the contemporary multiplayer strategy game *StarCraft* (Vinyals et al., 2019).

OT uses. The concept of reinforcement learning entered early on into the OT literature. In fact, a year after Samuel’s (1959) paper had been published, Clarkson and Simon (1960: 924–925) were already advocating for modeling learning processes in organizations using Samuel’s ideas. Reinforcement learning plays a central role in some of the works that defined the Carnegie tradition. Cyert and March (1963: 118) used the concept to explain a firm’s adaptive behavior: “Any decision rule that leads to a preferred state at one point is more likely to be used in the future than it was in the past.” Lave and March (1975: 247–339) used the Bush and Mosteller (1955) learning model to illuminate characteristics of adaptive behavior in organizations.

More recent work includes (a) Denrell and March (2001), which proposed the “hot stove effect” by which reinforcement learning can lead to more conservative decisions; (b) Denrell, Fang, and Levinthal (2004) and Fang and Levinthal (2009), which used Q-learning to shed light on how the effect of exploration in multistage decision problems differs from the effect in single-stage decision problems; (c) Rahmandad (2008), which used a reinforcement learning model to theorize about the effect of delays in the complexity of organizational learning; and (d) Puranam and Swamy (2016), which explored the nuances of learning processes in organizations in which joint actions determine organizational outcomes, which in turn feed back into the organizational learning process.

To conclude our discussion of the use of AI search approaches in OT (Approaches 1–4), we note that Winter’s (1987) seminal piece on “knowledge and competence as strategic assets” was critically influenced by all four AI approaches we have covered so far:

Both control theory and evolutionary theory invoke the notion of state description ... I have proposed the informal, looser and more flexible concept of a

³ TD-Gammon was also trained using self-play, but, where Samuel’s (1959) checkers program used feedback to change the weights of a polynomial, TD-Gammon used it to change the weights of a more complex neural net (see Approach 6 in the present paper).

heuristic frame—essentially, the control theory approach stripped down to a list of state descriptors and control. (Winter, 1987: 181)

Here is a clear demonstration that knowing AI can enrich our understanding of existing organizational theories and our ability to develop new ones.

AI IDEAS USED IN OT: REPRESENTATION APPROACHES

The emphasis of the previous set of approaches was that the key to producing intelligent behavior was to search a space of possible solutions. This could take the form of (a) searching for the value of a parameter so as to match an aspiration level (cybernetics and control theory, Approach 1); (b) searching a tree for a satisficing state (heuristic problem-solving, Approach 2); (c) searching a space of possible evolutionary designs (evolutionary computation, Approach 3), or (d) searching a space of possible policies to use in a dynamic program (reinforcement learning, Approach 4).

The premise of the next set of approaches is that AI can be achieved by picking the right representation. These approaches are very well characterized by Simon's (1969/1996: 132) assertion that "solving a problem simply means representing it so as to make the solution transparent." We will discuss three such approaches, differing in the type of representation they propose. *Expert systems* (Approach 5) focus on representations of logical rules, which have a good fit with mature knowledge domains for which well-defined explicit knowledge is available, such as the knowledge a doctor uses to diagnose diseases. (The optimism surrounding this approach spurred a countermovement that we also discuss.) *Connectionism* (Approach 6) focuses on statistical representations, which have a good fit with domains that rely on tacit knowledge, including perceptual tasks—such as recognizing faces or speech—for which high-quality explicit knowledge does not exist or is too difficult to acquire. *Bayesian networks* (Approach 7) focus on representing causal knowledge, which is particularly useful in domains characterized by uncertainty.

Approach 5: Expert Systems and Knowledge Representation

AI idea. "Expert systems" is an AI approach that relies on representing the knowledge of a domain as a large number of special-purpose rules that, when combined by an inferencing mechanism, can answer questions about the domain (Norvig, 1992: 461; see

also Feigenbaum, McCorduck, & Nii, 1988: 31–48). The idea that motivated the creation of expert systems was to create a general and flexible platform able to represent knowledge of different domains; that is, while a given expert system would be specialized, the approach used to create such systems aimed for generality (Norvig, 1992: 530). The expert system approach is thus in contrast to custom-made approaches, such as a chess program, that would be very difficult or impossible to adapt to a different domain. Expert systems achieve this flexibility by including two main components: a knowledge base and an inference engine. The first stores knowledge as if-then rules and the second combines that knowledge using logical rules. This architecture aims to mimic how an idealized expert uses accumulated knowledge to reason about his or her domain of expertise (Lindsay, Buchanan, Feigenbaum, & Lederberg, 1993).

An early example of an expert system is the MYCIN program for diagnosing bacterial infections (Shortliffe, Davis, Axline, Buchanan, Green, & Cohen, 1975). It used around 600 if-then rules, diagnosing an infection once enough evidence had been accumulated. A more recent example is the TurboTax software package, which uses myriad rules about the U.S. tax code to solve tax preparation problems.

Starting in the late 1960s, a number of successful expert systems were developed that could match the performance of human experts. These included Dendral (Buchanan et al., 1969), which could infer a molecular structure from the information provided by a mass spectrometer; MYCIN (Shortliffe et al., 1975)—mentioned above—which could diagnose blood infections; Macsyma (Moses, 1974), which could solve symbolic algebra problems; and PROSPECTOR (Duda, Hart, Nilsson, Reboh, Slocum, & Sutherland, 1977), which could recommend sites for mineral prospecting. The experience gained with these and other projects led to the creation of tools for creating expert systems, such as the programming language Prolog, and to a formalization of the process of extracting rules from experts, called "knowledge engineering" (Hayes-Roth, 1992).

Success also led to ambitious engineering projects such as Cyc (Lenat, Prakash, & Shepherd, 1985), whose aim was to produce a knowledge base that included all common-sense knowledge about how the world works, and Soar (Laird, Newell, & Rosenbloom, 1987), which aimed to provide a cognitive architecture that could produce human-level general intelligence.

In addition, new ideas arose that expanded the types of data and inferences that expert systems

could handle. Schank and Abelson (1977), for example, extended expert systems to operate not just on if-then expressions but also on “scripts”—knowledge structures that describe stereotypical situations such as going to a restaurant (including finding a table, choosing from the menu, and so on). Other structures developed to deal with stereotypical situations included frames (Minsky, 1975) and schemas (Bobrow & Norman, 1975). McCarthy (1980) and others (see Ginsberg, 1987) developed nonmonotonic logic to more easily deal with knowledge bases that include exceptions, contradictions, and default assumptions. This type of logic provided a workable solution to the “frame problem” (McCarthy & Hayes, 1969/1981), the notion that dealing with such situations using standard logic formalisms was computationally unwieldy.

The early successes of expert systems led to an “AI bubble” in the early 1980s, with the creation of numerous AI start-ups, many of which offered expert systems. The bubble was inflated in part by Japan’s Fifth Generation project, which injected close to \$1 billion (in today’s dollars) into the Japanese AI industry, and by the U.S. Defense Advanced Research Projects Agency’s reaction to that, which was to invest lavishly in AI projects through its Strategic Computing Initiative (Feigenbaum & McCorduck, 1983; Nilsson, 2010: 296). By the late 1980s, the boom had ended, due to the mismatch between the high expectations for expert systems and what they were actually able to deliver. The disappointments were emphasized by lingering academic doubts and criticisms about the potential of expert systems (Dreyfus, 1972; Dreyfus & Dreyfus, 1986). The end of the boom marks the beginning of the so-called “AI winter,” which lasted over 20 years, until advances in connectionist and machine learning approaches (Approach 6) raised new hopes for the field.

OT uses. The influence of expert systems on OT takes many forms. Interestingly, these influences have as much to do with the criticisms of expert systems as with the conceptual innovations that allowed for the existence of expert systems in the first place.

One manifestation of the influence of expert systems on OT is the efforts to codify parts of OT in the form of an expert system. In particular, Burton and Obel (2004) created OrgCon, which codifies much of what is known about organization design. Their expert system is structured in terms of if-then rules, such as, “If the organization is large, then decentralization should be high” (Burton & Obel, 2004: 18). It

includes, for example, rules about the role of interdependence (Thompson, 1967) and environmental uncertainty (Galbraith, 1973). Another manifestation is the adoption by Polos, Hannan, and Carroll (2002) and Hannan, Polos, and Carroll (2007) of nonmonotonic logic to express their theory about organizational forms and identities. Nonmonotonic logic allowed them to rigorously state their theory despite the fact that it deals with concepts that continually evolve and have fuzzy boundaries (e.g., like music and wine categories) and hence are ill suited for traditional formal tools like predicate logic or set theory.

Another set of OT uses of ideas originating with expert systems has to do with the limits of codified knowledge. Starting in the late 1980s, the failure of expert systems to match the high expectations they had created prompted several OT scholars to study the limits of codified knowledge. Prietula and Simon (1989: 120–121) highlighted the role of human experts and explained that expert systems usually fail to mimic them, as “[human] expertise is based on a deep knowledge of the problems that continually arise” in any particular kind of work. Disappointment with expert systems also influenced the development of the knowledge-based view, in which the process of “externalization”—turning tacit knowledge into explicit knowledge (Nonaka, 1994)—is described as imperfect. In this vein, Grant (1996) explained that “converting tacit knowledge into explicit knowledge ... inevitably involves substantial knowledge loss” and Kogut and Zander (1992: 387) cited Dreyfus and Dreyfus’s (1988) criticism of expert systems as supporting evidence for a similar view. Suchman (1987) and Orr (1996) undertook ethnographic studies that showed the limitations of an expert system created to support the work of copy machine technicians at Xerox. Drawing on these studies, Brown and Duguid (1991) theorized that organizational learning depends on “communities of practice” to facilitate creating and sharing knowledge that is situated and embedded in practice and which is quite distinct from the dry knowledge captured in an expert system.

The AI concept of scripts (Schank & Abelson, 1977) has been used in OT to theorize about routines and organizational adaptation. Nelson and Winter (1982: 79) built on the idea of scripts to theorize about the nature of skills and routines. Pentland (1995: 543) represented routines using a grammar that he characterized as “a more powerful generalization of the same basic idea of scripts.” Gioia and Poole (1984) saw much potential in investigating

scripts as a unit of analysis. They proposed a research agenda and methods for studying scripts in organizations, in response to which several papers studied how scripts change when firms face novel situations (see, e.g., Barley, 1986; Hargadon & Bechky, 2006; see also the multiple references to the concepts of script, frame, and schema—all AI concepts—in Walsh, 1995).

The idea of the frame problem (McCarthy & Hayes, 1969/1981) has also been used in OT to portray entrepreneurial action. Felin, Kauffman, Koppl, and Longo (2014) theorized that the frame problem implies that understanding entrepreneurial discovery as a process of boundedly rational search on a fitness landscape is an inadequate analogy. Following Dennett (1984), they interpreted the frame problem as implying that discovering all potential opportunities is computationally intractable, from which they drew the conclusion that it is incorrect to assume that there is any one landscape in which all these opportunities exist.

Approach 6: Connectionism and Machine Learning

AI idea. “Connectionism” is an AI approach that relies on artificial neural networks (also known as “neural nets”)—circuits whose connections are loosely patterned on those of the neurons in the brain (Russell & Norvig, 2020: 750; Sejnowski, 2018). Connectionism is part of a broader family of AI methods called “machine learning,” the goal of which is to build machines that improve automatically through experience (Jordan & Mitchell, 2015: 255). Some authors bundle machine learning and reinforcement learning (Approach 4) together. But a common way to differentiate between them is that, in reinforcement learning, the learning depends on trial and error—that is, on interaction with the environment—while, in machine learning, the computer infers relationships from data that are given—that is, learning depends on experience rather than on experimentation (Sutton & Barto, 2018: 2).

The main challenge that machine learning tries to address is how to use vast amounts of data to make accurate predictions and classifications. To do so, machine-learning systems need to explore “a large space of candidate programs, guided by training experience, to find a program that optimizes the performance metric” (Jordan & Mitchell, 2015: 255). Two examples of tasks will illustrate two types of machine learning. The first task is as follows:

Given the buying histories of a company’s customers, group them into similar types.

This is an example of “unsupervised learning,” as the data set only includes independent variables. The second task is:

Given the buying histories of a company’s customers and their ages, predict the age of new customers from the items in their shopping carts.

This is an example of “supervised learning,” as the data set contains both dependent and independent variables, as if a teacher or supervisor had labeled a training set of examples.

McCulloch and Pitts (1943) proposed the first mathematical model of how biological neural networks work. Simplifying some historical details, their model is currently understood as follows:


- Each neuron computes a weighted sum of its inputs and returns a scaled number that can be used as a result or as an input by other neurons (see Panel A in Figure 2).
- A neural network is an interconnected arrangement of such neurons (see Panel B in Figure 2).
- Most neural nets are hierarchical: an initial layer of neurons receives inputs from the external world (i.e., something outside the network itself) that are processed by subsequent layers and a final layer of neurons returns a result.
- The information passed from one layer to the next can be thought of as successively refined, higher-level representations of the original incoming information. For example, if the incoming information is all the pixels of an image and the last output of the net is a decision on whether or not the image contains a human, intermediate layers could be encoding successively higher-level concepts such as curves, blobs, and face-like features. The intermediate layers are sometimes known as “hidden layers.”
- Part of the appeal of neural networks is that the inputs and outputs can be anything that can be turned into numbers, such as the letters in a sentence, the pixels in a video, or a price–time series.

McCulloch and Pitts’s (1943) model was speculative and today it is well known that this is not an accurate representation of actual biological neural networks.⁴ Nevertheless, their model was intriguing


⁴ Artificial neural networks are not “biologically faithful,” as they do not account for myriad biological aspects such as neuron types, neurotransmitters, connectivity patterns, and brain structures (Crick, 1989). For details, see Hasson, Nastase, and Goldstein (2020: 417) and references therein.

FIGURE 2
Illustration of an Artificial Neuron and a Neural Net

A An artificial neuron



B An artificial neural net



and spurred much interest in practical applications. Decades later, this approach made it possible to solve several AI challenges better than any other AI approach could do and sometimes even better than humans.

Minsky (1954), Rosenblatt (1958), and others implemented neural networks with a single layer of neurons, called “perceptrons.” Minsky and Papert’s (1969) book then demonstrated that perceptrons had fundamental limitations in the types of functions they could compute, provoking a decline of interest in neural nets. Consequently, AI was dominated, in the 1970s and 1980s, by symbolic approaches—particularly, heuristic problem-solving and expert systems. In 1982, however, Hopfield (1982) developed a neural net that could be used as associative memory. That is, it retrieves the memory that is most similar to the stimulus with which it is presented. And, in 1986, neural nets took a big step forward when Rumelhart et al. (1986) published their “backpropagation” algorithm

for determining the weights of arbitrarily deep neural networks. Up to that point, there had been no workable method to do this, but now it was possible to overcome the limitations of single-layer neural networks and a door was opened to a great deal of exploration of various neural network architectures and uses.

Several other mathematical advances also furthered the development of neural networks. The demonstration that neural networks could approximate arbitrarily well any mathematical function (Cybenko, 1989; Hornik, Stinchcombe, & White, 1989) meant that, at least in theory, neural nets could be used to perform any observable cognitive function, since all sense data and actions can be encoded numerically. The derivation of the bias–variance decomposition (Geman et al., 1992) helped establish the optimal complexity of a neural network as a function of the data available for “training”—that is, for estimating the parameters or “weights” of the neural net. The use of “decision boundaries” made it

possible to visually compare the capabilities of different machine-learning approaches (Nilsson, 1965: 4–5; Vapnik, 1995).

Enabled by massive improvements in computer power and affordability, AI researchers discovered that neural nets with many layers—in some cases, thousands of layers; hence the name “deep neural nets” or “deep learning”—could surpass other AI methods in a number of tasks, including scene recognition, face recognition, transcription of speech, and translation (see Jordan & Mitchell, 2015, for an overview).

Combining neural nets with heuristic search (Approach 2) and reinforcement learning (Approach 4) allowed computers to defeat, for the first time, the human champions in the games of backgammon (Tesauro, 1995), Go (Silver et al., 2016), and *StarCraft* (Vinyals et al., 2019). In such systems, the neural net provides a scoring function used to evaluate possible configurations. The weights in the neural net are optimized via reinforcement learning from self-play and from supervised learning on games played by human experts.

OT uses. Neural nets have entered OT in two main ways: as an overarching model of information processing in organizations and as a way to think about learning and interpretive processes in organizations.

Inspired by McCulloch and Pitts (1943) and by cybernetics (Approach 1), Beer (1972) saw the organization as a type of neural net in which information was processed as it flowed up through the organization. From this observation, Beer concluded that organizations could be improved by rewiring them: adding sensors, routing the right information to the right people, and speeding up communications. His most ambitious application of these ideas was to create an information system to run the entire economy of Chile during the 1970–1973 socialist administration of Salvador Allende. Although the system was designed and partially implemented, there was little time for it to be used given the coup that abruptly ended Allende’s government (see Medina, 2011, for details about this information system).

The concept of neural nets was also used in several ways to model an organization’s learning and interpretive processes. Marchiori and Warglien (2008) modeled interactive learning in repeated games using neural nets to represent the decision-makers. Here, individual learning corresponds to updating the weights in a neural net representing each individual. In many cases, their model provided more accurate predictions of human behavior than competing economic models did. Gavetti and Warglien (2015) used

Hopfield’s (1982) neural model of associative memory to represent individuals’ interpretive processes. They then connected these individuals in a social network to study whether the extent to which paying attention to others’ interpretations leads to more accurate group-level decisions. They found that the benefit of paying attention to others follows an inverted-U relationship: paying too little attention to others misses valuable insights, while paying too much attention to others results in conformism and neglects the actual problem. Csaszar and Ostler (2020) used the bias–variance decomposition to theorize about the optimal complexity of the representations used by organizations as a function of managers’ level of experience and the complexity and uncertainty of the environment. Their theory provides a way to reconcile conflicting views on the debate of simple versus complex representations (i.e., Bingham & Eisenhardt, 2011, vs. Weick, 1979: 261) by characterizing situations under which representations of low, medium, and high complexity are better able to make accurate predictions about the environment.

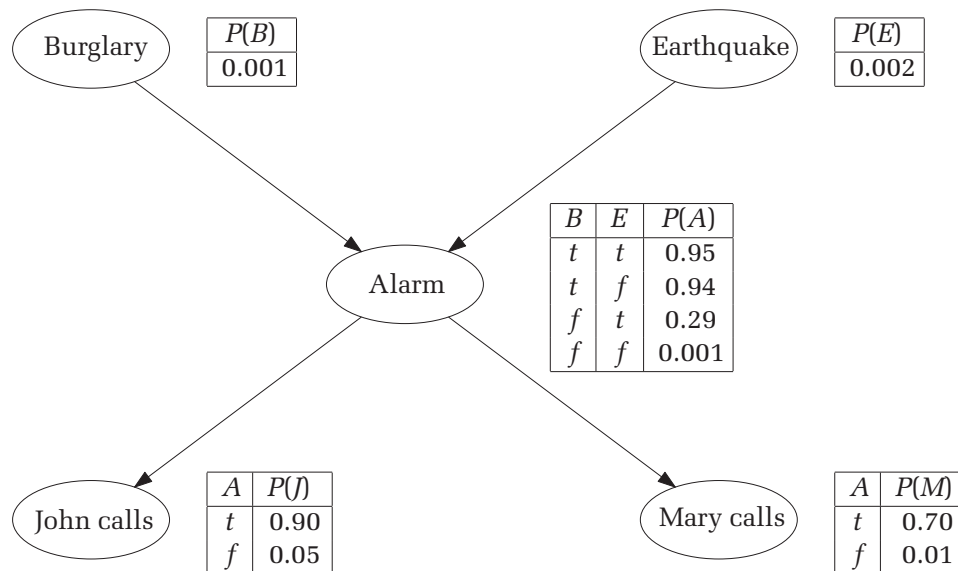
Although the scope of the present paper is to cover AI analogies in OT, it is worth noting that an important effect of the current wave of machine learning on OT is a methodological one: to expand the type and amount of data available to researchers. For instance, researchers are using machine-learning methods to categorize texts, code videos, and discover patterns in the data they have collected (see Choudhury, Allen, & Endres, 2021 for a review).

Approach 7: Bayesian Networks

AI idea. “Bayesian networks” (also known as “belief nets” and “graphical models”) are an approach to AI that allows programs to reason probabilistically about causes and effects (Nilsson, 2010: 381). Dealing explicitly with probabilistic situations contrasts with previous approaches, such as heuristic problem-solving and expert systems (Approaches 2 and 5), that used logic to deal with true/false propositions but did not have a sound way to represent probabilistic relationships and make probabilistic inferences.

There had been previous attempts in AI to extend logical approaches to better capture uncertainty, such as Dempster–Shafer logic and fuzzy logic, but these had been found to be unsound (Koller & Friedman, 2009: 13). Bayesian networks build on the insight that following the laws of probability is the only way to behave rationally in the face of uncertainty (de Finetti, 1937/1980, Ramsey, 1931). To

FIGURE 3
Illustration of a Bayesian Network



accomplish this, Bayesian networks provide a language to represent and interconnect arbitrarily detailed probabilistic information and an inference mechanism that uses the Bayes theorem to update arbitrarily large and interconnected webs of beliefs.

A canonical example of a problem calling for the use of a Bayesian network is illustrated in Figure 3:

Given the causal structure and probabilities linking a home alarm going off to other events, determine the probability that the home was actually burglarized if, say, the alarm goes off, Mary called, and there was an earthquake.

A more sophisticated use of Bayesian networks is to infer the causal structure from data. For example, Friedman, Linial, Nachman, and Pe'er (2000) inferred complex networks of gene interactions from observing gene expression data. The idea was to recover the causal structure of gene interactions (i.e., what genes affect what genes) from examining statistical properties of dependence and conditional independence in the data.

The idea of using graphs to encode causal information can be traced back to a paper by Wright (1921) on genetics (Pearl, 2000: 26).⁵ Pearl (1988) popularized the notation used in Figure 3 (called a “Bayesian

network”) and developed efficient algorithms to make inferences on these networks. The key insight of these algorithms is to use the dependence structure of the Bayesian network to avoid the combinatorial explosion that would otherwise occur when computing the problem’s joint probability distribution.⁶

Cooper and Herskovits (1992) developed the first algorithms to “learn” or infer Bayesian networks from data. More recent developments include the use of Bayesian networks in cognitive science as a benchmark of rational decision-making against which to compare cognitive models and experimental data (Oaksford & Chater, 2007). Cognitive science has also used Bayesian networks to model processes of discovery (Spirtes et al., 2000) and concept learning (Tenenbaum & Griffiths, 2001).

OT uses. Most of the OT uses of Bayesian networks have to do with modeling the strategic decision-making process. Durand and Vaara (2009), for example, proposed a research agenda based on

⁶ For instance, if, in Figure 3, one knows that the alarm went off, knowing whether there was an earthquake or not does not change the probability that Mary will call. This is an illustration of the “d-separation” criterion (Pearl, 2000: 16), one of the ways by which Bayesian network algorithms make inferences without having to evaluate all the possible combinations of contingency values—something that would be prohibitively costly except in relation to the simplest problems.

⁵ Coincidentally, Wright is also the originator of the rugged landscape analogy commonly used in OT.

understanding the relationship between resources and performance in terms of Bayesian networks. They saw Bayesian networks as a good description of the processes driving firm performance and saw the role of the strategist as understanding that causal structure and using it to make well-informed interventions.

In line with Durand and Vaara (2009), Ryall (2009) used Bayesian networks to develop a model of causal ambiguity (Lippman & Rumelt, 1982) that represented (a) the causal structure of a firm's activity system and (b) managers' beliefs about that activity system. Using these concepts, Ryall (2009) theorized about two types of causal ambiguity—intrinsic and subjective—stemming from the uncertainty inherent in the real and the perceived Bayesian networks, respectively.

The agenda of embracing Bayesian networks in strategy has also affected how strategy is taught and studied. The textbook by Ryall and Bramson (2014) teaches Bayesian networks to MBA students on the premise that strategy consists of making rational decisions under conditions of uncertainty and that Bayesian networks are the best embodiment of the principles of rational decision-making under such conditions. In terms of research methods, Bettis and Blettner (2020) built on Spirtes et al. (2000) and called for studying causality using Bayesian networks rather than regressions. The argument is that regressions essentially capture correlations, whereas Bayesian networks are better suited to describe and estimate the causal structure that gives rise to the observed data.

AI IDEAS USED IN OT: AGGREGATION APPROACHES

We now move on to approaches to AI that rely on aggregation. These differ from the previous approaches in that, here, intelligent behavior emerges from combining the actions of several subsystems or agents to create a desired aggregate behavior. There is no restriction on the types of subsystem that can be used. They could be other AI systems (including any of the approaches we have covered so far), but they could also be simpler decision rules, such as hand-coded rules and statistical methods (e.g., linear regressions). The key to aggregation approaches is that—like the U.S. motto "*E pluribus unum*"—out of many parts, a superior aggregate behavior is achieved.

Approach 8: Cellular Automata and Emergence

AI idea. "Cellular automata" is an approach to AI that aims to produce intelligent behavior by

imitating the functioning of an idealized biological cellular tissue. The "cells" in this "tissue" are described by a state that evolves based on rules that depend on the state of neighboring cells (Floreano & Mattiussi, 2008: 101). More formally, a cellular automaton has two components: (1) a grid of N identical cells, each with an identical pattern of communication with neighboring cells, and (2) a transition rule, which describes how a cell changes from one time period to the next (Mitchell, 1998: 96).


The challenge the cellular automata approach addresses is how to make intelligent behavior emerge from very simple rules. Note that this is in direct contrast with all the other approaches we have seen so far, which use more complex structures, such as a large knowledge base or myriad weights in a neural net, and whose behavior is not emergent—that is, their behavior is directly programmed rather than induced by the system's intrinsic dynamical behavior (Hanson, 2009). Arguably, this approach is representative of the types of structure that social or biological evolution can create.

Research on cellular automata and emergence has been intellectually fruitful, but so far has not produced technological breakthroughs like those attributable to the other AI approaches. For this reason, cellular automata is not usually included in introductory AI textbooks (an exception is Luger, 2005), but only appears in more specialized ones (e.g., Adami, 1998; Flake, 1998; Floreano & Mattiussi, 2008).

Perhaps the most famous cellular automaton is John Conway's Game of Life, which creates extremely rich, life-like behavior based on just four rules determining when a cell becomes dead or alive (Gardner, 1970). Watching an animation of this automaton resembles seeing a world teeming with microscopic life; Figure 4 shows one snapshot of this automaton in action. Another famous cellular automaton is Wolfram's (1983) Rule 30, a one-dimensional cellular automaton that produces unpredictable complex behavior. Wolfram (2020) and others (Schiff, 2008: 181; 't Hooft, 2016) have even speculated that the universe is a cellular automaton not too different from Rule 30.

Von Neumann pioneered the idea of cellular automata in his proposal for a general mathematical model of a self-reproducing machine, called the Universal Constructor (von Neumann & Burks, 1966). Burks's (1970) book helped define cellular automata as a multidisciplinary academic field. Conway's Game of Life created broad interest in cellular automata, helped by its wide diffusion in *Scientific American* (Gardner, 1970). Varela et al. (1974) used

FIGURE 4
Snapshot of a Cellular Automaton



cellular automata to develop the concept of “autopoiesis,” the capacity particular to living systems of reproducing and self-healing. Langton (1986) built on von Neumann’s idea of self-replication to pioneer the field of artificial life, which used cellular automata to understand biological phenomena such as biochemical reactions, ant colonies, and cellular replication. Bak, Tang, and Wiesenfeld (1987) used cellular automata to describe the phenomenon of “self-organized criticality”; that is, systems that are able to maintain a characteristic structure via an emergent self-correcting mechanism. For example, avalanches keep the shape of all sandpiles similar regardless of how much sand is in the pile. Crutchfield and Mitchell (1995) combined cellular automata with genetic algorithms (Approach 3) to evolve cellular automata that produce a desired behavior. The long-awaited publication of Wolfram’s (2004) book, *A New Kind of Science*, renewed the visibility of cellular automata.

OT uses. In OT, cellular automata have been used to represent social processes that exhibit spatial distribution and local connectivity, such as processes of segregation, competition, collaboration, and diffusion. For example, Schelling’s (1969, 1971)

celebrated studies about segregation model a city as a cellular automaton in which individuals change location depending on the proportion of their neighbors that they consider to be of the same type. Disturbingly, the model shows that, even if that desired proportion of similar neighbors is small, cities are likely to devolve toward segregated neighborhoods. That is, even if individuals don’t seek what they would consider segregation, their individual choices may collectively produce it.

Nowak and May (1992) extended the research on the evolution of cooperation (Axelrod, 1984) to explore the societal conditions that make cooperation or competition more likely to emerge. In their model, individuals play an iterated prisoners’ dilemma with their neighbors, with the winning strategy becoming more likely to be diffused.

Lomi and Larsen (1996) used a cellular automaton to model the population dynamics of organizations. They show that this simple model can replicate classical findings of the organizational ecology literature (Hannan & Freeman, 1989). Their model also suggests that previous results on density dependence are very sensitive to changes in the rules of local interaction; that is, to the cellular automaton’s transition rules.

Lustick (2000) used cellular automata to theorize about how identities diffuse. In this model, the activation of an individual's identity depends on the neighbors' identities. (For introductions to the use of cellular automata in the social sciences, see Hegselmann, 1996; Nowak & Lewenstein, 1996; and Schiff, 2008: 123–184)

Finally, it is worth noting some more distant influences of cellular automata on OT. Much research in OT has built on the complex adaptive systems literature, which develops the themes of emergence, chaos, and complexity by using cellular automata as one of its main methodological tools (Holland & Miller, 1991; Miller & Page, 2007; Mitchell, 2009; Waldrop, 1992). Examples of OT literature building on this work include Anderson (1999) and Allen, Maguire, and McKelvey's (2011) edited volume. Eisenhardt and Piezunka (2011), included in Allen et al.'s (2011) volume, highlighted the strategy literature's links with the complex adaptive systems literature.⁷

Luhmann's (1990) theory of autopoietic social systems extended Varela et al.'s (1974) ideas to social—not just biological—systems. In this version, social systems reproduce themselves on the basis of communication (Hernes & Bakken, 2003; Seidl, 2004). All told, it is fair to say that, without cellular automata, several models about organizational dynamics and the OT literatures on autopoiesis and complex adaptive systems would probably not exist.

Approach 9: Distributed AI

AI idea. “Distributed AI” (also known as “multiagent systems”) is an approach to AI that aims to achieve intelligent behavior by using multiple interactive intelligent agents. These agents are autonomous programs capable of interacting with other agents via “social” behaviors such as cooperation, coordination, and negotiation (Weiss, 2013: xxxv; Wooldridge, 2009: xiii). Distributed AI differs from cellular automata in that distributed AI is much less constrained in the range of behaviors and communication structures available to agents. Under this approach, agents are usually not controlled by the extremely simple rules guiding the evolution of a cellular automaton. Instead, they can pursue goals,

sense the environment, communicate with others, and make choices. Also, distributed AI allows for a more flexible communication structure that is not restricted by the regular grid of cellular automata.

The challenge that distributed AI is trying to address is how to deal with situations in which each agent has access to incomplete information or has incomplete capabilities and the agents as a group cannot be centrally controlled (Sycara, 1998). A current example is a system in which self-driving cars share information about road conditions and about the other cars on the road, allowing each car to drive more safely and rapidly by avoiding accidents, coordinating maneuvers, and choosing less-congested roads. Another example would be a team of robots sharing a task, such as those competing in the annual RoboCup soccer competition (Kitano, Asada, Kuniyoshi, Noda, Osawa, & Matsubara, 1997).

The earliest distributed AI program was Selfridge's (1959) Pandemonium, which could recognize written characters using multiple agents running in parallel. The agents would detect specific features of letters, such as vertical lines, horizontal lines, and circles, and would relay their opinion to agents specialized on higher-level tasks such as recognizing the overall shape of a letter. Similarly, Erman, Hayes-Roth, Lesser, and Reddy (1980) showed how speech recognition could be drastically improved by having multiple simple agents working at different levels of analysis—for example, the letter, word, and phrase levels—and having those agents share information through a shared “blackboard” system.

Minsky's (1986) “society of mind” concept proposed that the human mind works in a way similar to that of Selfridge's system; that is, with several agents working in parallel, each serving a different human need. According to this theory, people only become consciously aware of those agents that “shout” loud enough.

Winograd and Flores's (1986) philosophical text on the nature of cognition and distributed computation highlights the problems of coordinating actions and achieving shared understanding in distributed systems. To examine these problems, they modeled communication using diagrams that describe the possible states and transitions in a dialogue that coordinates an action. An early commercial system implementing these ideas was described by Flores, Graves, Hartfield, and Winograd (1988).⁸

⁷ An example of the link between the OT literature on complex adaptive systems and cellular automata is the concept of “edge of chaos” (used in OT; e.g., by Brown & Eisenhardt, 1998, and Carroll & Burton, 2000) and which emerges from Langton's (1986) characterization of the class of cellular automata that exhibit chaotic behavior.

⁸ Interestingly, it was Flores who hired Beer in 1971 to work on Cybersyn, the system intended to run the Chilean economy mentioned under the OT uses of Approach 6.

The growing availability of computer networks in the 1980s increased the applicability of distributed AI, as answers to questions like these became consequential: How can an automated agent best represent the interests of a user in an online negotiation with other automated agents (Rosenschein & Genesereth, 1985)? In what ways should humans and computers interact in the office of the future (Hewitt, 1988)? In the late 1980s, the first books on distributed AI were published (Bond & Gasser, 1988; Huberman, 1988). Since then, the field has continued to grow, often by building on theories originally developed to understand human behavior. A recent trend has been to develop distributed AI systems that incorporate ideas from economics, such as game theory, mechanism design, and auction theory (see, e.g., the introductory textbooks by Wooldridge, 2009, and Weiss, 2013, and, in particular, the more specialized volume by Shoham & Leyton-Brown, 2009).

OT uses. One use of the concept of distributed AI in the field of OT has been computational organization theory, which Carley and Gasser (1999: 323) described as distributed AI models informed by empirical knowledge from organization science. They also described several tools from the distributed AI literature that could serve organizational researchers. The article by Carley and Newell (1994) suggested how to characterize the types of agents and situations that this approach may study, and the edited volumes by Carley & Prietula (1994) and Prietula, Carley, and Gasser (1998) provided many examples of this approach.

Another use of distributed AI in OT is representing organizational processes. Malone and Crowston (1991: 6), for example, proposed developing a “coordination theory”—based on insights from distributed AI, economics, OT, and biology—to understand organizations. An ambitious outcome of this research agenda was the book by Malone et al. (2003), which presented a technical language to describe organizational processes. The goal of this language was to organize all business knowledge and to turn organization design into a discipline closer to engineering, which they suggested would lead to improvements in the quality, predictability, and reusability of organization designs. Their proposed language built on languages used in distributed AI, such as Winograd and Flores’s (1986) transition diagrams (for other examples, see Crowston, 1992).

Inspired by distributed AI ideas, Wegner and his collaborators introduced the theory of transactive memory systems (Wegner, 1987; Wegner, Giuliano, & Hertel, 1985) in the context of intimate dyads:

couples who have a shared understanding of “who knows what” and who have processes for encoding, storing, and retrieving such shared information. The concept was later extended to groups and organizations (e.g., Liang, Moreland, & Argote, 1995), and has led to a rich empirical literature (see Ren & Argote, 2011, for a review). At a conceptual level, transactive memory systems are appealing because they explain much of organizational cognition without resorting to the dubious concept of a “group mind.” Instead, transactive memory systems are faithful to the cognitive limitations of the individual actors and to the distributed nature of work in the same way that distributed AI is.

Wegner (1995: 320–321) pointed to the distributed AI roots of the concept of transactive memory systems by discussing how these systems can best be understood by imagining building software that finds information in a computer network—a quintessential distributed AI problem (see, e.g., Hewitt, 1977). He cited foundational distributed AI works (Braitenberg, 1984; Minsky, 1986) to justify the value of this thought experiment.

Yet another OT concept with AI roots is boundary objects. Carlile (2002) introduced the concept of a boundary object to OT. In this context, a “boundary object” is an artifact that enables collaboration, such as a shared pad of paper or a 3D model. But the idea comes from a chapter in one of the foundational books on distributed AI, in which Star (1989: 51) proposed the concept of the boundary object as a solution to the problem of sharing information in the “blackboard” systems popular in the distributed AI literature at that time. Star herself is also a good illustration of the two-way connection between AI and OT, as she used organizational examples to motivate different classes of boundary objects. Boundary objects continue to be an active research area in OT (see, e.g., Bechky, 2003; Seidel & O’Mahony, 2014; Zuzul, 2019).

Approach 10: Ensemble Methods

AI idea. “Ensemble methods” are an approach to AI that relies on combining predictions from different programs (called “learners” in this literature) that work on the same or very similar data. Alternative names for this approach are “committee-based learning” and “multiple classifier systems” (Zhou, 2012: 15). Ensemble methods differ from distributed AI in that ensemble methods combine similar programs performing a similar task, while distributed AI combines programs that are each specialized in a

given task (such as the players in robot soccer or the various subsystems in a driverless car). In other words, ensemble methods rely on redundancy while distributed AI relies on division of labor.

The challenge this approach is trying to address is to avoid the weaknesses of any single learner by combining several learners. Combining learners is appealing because there is no infallible way to find the best learner—a result known as the “no free lunch” theorem (Wolpert, 1996, 2013). The usual way to deal with this difficulty is to try out several learners and pick the one that performs the best. Ensemble methods offer an alternative to that approach: rather than relying on just one learner, find a way to combine learners so as to collectively outperform any single one. (Or, in machine learning lingo, combine several weak learners to create one strong one.)

A canonical example of ensemble methods is the “wisdom of the crowd” effect (Surowiecki, 2004), first described by Galton (1907/1949), who observed that the median guess in a competition to find out the weight of an ox was more accurate than the predictions of almost all the participants. A more recent example is the Pragmatic Chaos program, which won a \$1 million prize from Netflix for improving the accuracy of its predictions of movie recommendations by more than 10%. This program worked by combining predictions from hundreds of learners (Koren, 2009).

The idea of improving predictive accuracy by using ensembles has a long history and builds on three foundational ideas. The first is Bernoulli’s (1713/2006) law of large numbers, which states that averaging many independent samples will converge to the true population average as the number of samples increases. It follows that combining several independent learners should result in more accurate predictions.

The second foundation of ensemble methods is Condorcet’s (1785/1994) jury theorem. Conceptually similar to the law of large numbers, but in the context of voting, it states that majority voting conducted by individuals who have a better-than-chance probability of picking the best alternative converges (in probability) to picking the best alternative as the number of individuals increases. (Condorcet came up with this result in the lead-up to the French Revolution, as a way to understand whether democracy could work better than monarchy.)

The third foundation is reliability theory; in particular, seminal work on how to create reliable electronic circuits out of notoriously unreliable parts

such as vacuum tubes (Moore & Shannon, 1956/1993; von Neumann, 1956). The idea is to be able to use the structure of a circuit to compute the probability that the circuit will produce an accurate result. This knowledge can then be used to create circuits with enough redundancy so that they work as intended with an arbitrarily high probability.

Key milestones in the modern history of ensemble methods include the proposal of various ideas on how to create ensembles. One such idea is to create ensembles by using different families of learners (e.g., neural nets, regressions, and clustering) or different learners within a given family (e.g., neural nets with different numbers of nodes and layers). Hansen and Salamon (1990) noted that, for an ensemble to produce correct results, it is important that its learners make uncorrelated errors. This is a strong argument for using a diverse set of learners.

A second idea was to train similar learners on different subsets of the data (i.e., different “rows” of the data set). This is a cheap way to create a diverse *set* of learners out of only one *type* of learner (say, neural nets). The most common way of doing this is “bagging” (Breiman, 1996), which works by training learners on subsets of the original data set generated by sampling with replacement.⁹

A third idea was to train similar learners on different attributes of the data (i.e., different “columns” of the data set). The method of “random forests” (Ho, 1995) introduced this idea by training decision trees on the same data set, but picking at random the variable used at each branching point of the tree.

A fourth idea was to train a sequence of learners such that successive learners try to correct the errors made by previous learners. Unlike the first three ideas, which use multiple learners working in parallel, this idea works by placing the “experts” in sequence. Schapire (1990) introduced this approach, which he called “boosting.” A common version of the boosting idea is AdaBoost (Freund & Schapire, 1996), in which each successive learner is trained on a sample that assigns more weight to the predictions that previous learners got wrong.

Following the law of large numbers and Condorcet’s jury theorem, most ensemble methods use simple averaging or voting (depending on whether the task is, respectively, prediction or classification). It is also possible to use weighted versions of averaging and voting. An important result that has been

⁹ The name “bagging” comes from bootstrap aggregating, as it operates similarly to bootstrapping in statistics (Efron & Tibshirani, 1993).

rediscovered many times is that the optimal weights for uncorrelated learners are their log odds of being correct (i.e., $w_i^* = \log(a_i/(1-a_i))$, where a_i is the accuracy of learner i). Pierce (1961) seems to have been the first to discover this result (Kuncheva, 2014: 125). The theory of Bayesian model averaging provides a general way of computing optimal weights when the learners are correlated.

Dietterich (2000) provided three explanations for the success of ensemble methods. The *statistical* reason is that, by using multiple learners, the ensemble decreases the risk of picking a bad one. The *computational* reason is that estimating all the parameters in a learner is a computationally complex process prone to getting “stuck” at a local peak, whereas combining multiple learners can better approximate the optimal parameters. The *representational* reason is that most learners are unlikely to have the correct representation of the problem—they may, for example, be missing variables or ways of interacting them—whereas a combination of learners will be more able to represent a complex problem correctly.

OT uses. Several OT papers have used ideas from the literature on ensemble methods to conceptualize how individuals and organizations make decisions. One group of uses has to do with information aggregation in organizations. Csaszar and Eggers (2013) modeled small groups tasked with screening projects—specifically, selecting between good and bad projects. The authors conceptualized these groups as ensembles of individuals and compared the effect of three aggregation rules: voting, averaging, and delegation. They showed, for example, that majority voting performs better than the other rules in most situations. Csaszar (2013) developed a model of organizations as ensembles and examined how the omission and commission errors made by an organization depend on its aggregation structure. One result stemming from this work is that some aggregation structures are always dominated by others and, hence, should never be chosen. Csaszar (2012) tested the effect of aggregation structure on the probability of making omission and commission errors (a relationship first theorized by Sah and Stiglitz, 1986). Using mutual funds as an empirical context, Csaszar (2012) showed that funds employing unanimous decision-making make fewer commission errors but more omission errors (i.e., pursue fewer failed investments but miss more good investments) than funds that do not require unanimity (and vice versa).

Ensemble methods have also been used to investigate the reliability of organizational structures. Christensen and Knudsen (2010), for example, built on

Moore and Shannon (1956/1993) and Sah and Stiglitz (1986) to examine how an organization’s structure affects its reliability—that is, the probability that it will pick good alternatives. Like Moore and Shannon, Christensen and Knudsen (2010) described structures in terms of circuits, which can connect individuals either sequentially or in parallel. Their paper illustrates how to design organizations that achieve a desired level of reliability. In a similar vein, Knudsen and Levinthal (2007) developed a model that shows how organization structure affects reliability, which in turn affects firms’ ability to explore new alternatives.

Another group of OT uses relates to the wisdom of the crowd and idea selection. Grushka-Cockayne et al. (2017) modeled crowds as random forests. That is, although members may observe similar data, they end up with different models because learning is an idiosyncratic process. Csaszar (2019) modeled crowds as an ensemble using majority voting, and studied how the probability of making a correct decision depends on crowd size, the accuracy distribution of the crowd, and the firm’s ability to recruit accurate individuals to be members of that crowd. One finding of this research is that, under relatively common conditions, increasing the size of the crowd may actually decrease the accuracy of predictions. Because it is difficult to reliably assess the accuracy of individuals in a crowd, Graefe, Küchenhoff, Stierle, and Riedl (2015) showed that, in realistic cases, it is typically preferable to use equal weights than more sophisticated methods. Page (2007) summarized much of the literature on why groups can outperform individuals, using the arguments that explain the success of ensemble methods.

Ideas about ensemble methods have also been used to study the relationship between cognitive diversity and decision quality. Page (2018: 30) used ensembles to conceptualize individual-level decisions. Using the arguments about the superiority of ensembles, he proposed that individuals who use multiple models to understand a phenomenon make better decisions than those who rely on one model. An early application of ensembles to understand individual cognition is Arthur’s (1994) famous El Farol problem, which modeled individuals as holding multiple hypotheses—that is, an ensemble of theories about how the world works.

DISCUSSION

We have shown that OT has borrowed many ideas from AI and that the breadth of the borrowing is staggering, ranging from cognitive diversity to

organizational reliability to exploration, from aspiration levels to organizational learning to requisite variety, and from scripts to transactive memory to the wisdom of the crowd, just to name a few.

This wide range is a result of the fundamental connection between AI and OT. Both aim to produce intelligent behavior—one with computer chips, the other mostly with humans. March (1999) said that OT is about the “pursuit of organizational intelligence,” which is not too different from what AI is. For this reason, fundamental aspects of intelligence—like search, representation, and aggregation—matter to both AI and OT. And because the connection is at such a fundamental level, the parallels between AI and OT are ubiquitous; indeed, we have cited over 100 OT papers that critically depend on AI ideas. This deep connection also explains why the borrowing has not only been from AI to OT, but also the other way around (as outlined at the end of the Background section and elaborated at different points in the paper).

It is surprising that these deep and diverse linkages between AI and OT are not usually acknowledged, unlike OT’s linkages to economics, psychology, sociology, and evolutionary biology. (See, e.g., the limited role that AI plays in classic introductions to the OT field, such as Morgan, 2006, and Scott & Davis, 2007.) Overall, the impact of AI on OT may be one of OT’s best-guarded secrets.

We believe that understanding AI’s influence on organizational theories is important not just for the sake of intellectual honesty and curiosity but also (a) to improve scholars’ ability to understand and create OT theories that build on AI and (b) to increase the repertoire of ideas with which to understand organizations. Note that the second reason entails the benefits of cognitive diversity mentioned in the context of ensembles (Approach 10).

The bulk of this paper has been backward looking, as the goal has been to review and explain the existing connections between AI and OT. In this final section, however, we take a forward-looking view to discuss three topics: first, how some AI ideas have evolved after being adopted by OT and thus may be ripe for new borrowing; second, how AI may continue to inform OT research (in particular, we point out possible AI–OT analogies that have not been explored so far); and, third, how AI itself may affect organizations.

How AI Ideas Have Evolved After Being Adopted by OT

The bridges between AI and OT have usually been established with the idea that was most popular in AI

at the time. For example, ideas about heuristics in OT (Simon, 1955) were imported around the time that the heuristics approach was making big strides in AI—for example, with the Logic Theorist and the General Problem Solver. Similarly, the OT idea of understanding organizational adaptation as search (Levinthal, 1997) was developed at a time when search was a popular AI technique. In fact, it was in 1997 that IBM’s Deep Blue used a search-based algorithm to beat the world chess champion, Gary Kasparov.

But, after OT imports an idea from AI, AI continues to move on while its OT “copy” often remains frozen. Looking at how AI approaches have diverged since being imported into OT can help us obtain a more realistic view of AI than the view prevalent at the time of the borrowing, appreciate the strengths of newer methods, and evaluate whether it makes sense to rethink some of the AI-based analogies OT has used. Three “post-borrowing divergences” that are instructive in these respects are the divergence between expectations and reality, the divergence between human and computer capacities, and the divergence among different factions within the AI community.

The divergence between AI expectations and actual achievements is well illustrated by the history of expert systems. In the 1980s, there was much anticipation about the ability of expert systems to replace human decision-makers. A string of unmet promises led to an “AI winter,” which lasted until the last decade, when new algorithms and computing power allowed connectionist ideas (Approach 6) to achieve human-level performance for the first time on several recognition tasks. Being aware of AI failures is useful to OT scholars, as it tempers predictions about impending AI scenarios.

The divergence between human and computer capacities stems from the fact that, while human capacity has remained fixed over the history of AI, computer capacity has increased exponentially. This has made the AI approaches that are more computing intensive become more powerful over time, giving them an edge over the approaches that depend less on computing power and more on human knowledge. Sutton (2019: para. 1) called this the “biggest lesson that can be read from 70 years of AI research.” For instance, the initial theories about computer chess imagined the solution was going to be a large collection of heuristics (the initial estimate by Newell & Simon, 1976: 125 was that an expert-level chess program would comprise about 50,000 such heuristics). In contrast, the program that ended up beating humans at chess contained few heuristics

and instead relied heavily on brute force search (Deep Blue analyzed over 100 million positions per second; Campbell et al., 2002). Similarly, while the early approaches to image recognition relied on hand-crafted rules, all the approaches that currently dominate this area are based on neural nets trained on massive data sets (Jordan & Mitchell, 2015). One implication is that the AI approaches that depend more on computing power (e.g., connectionism, ensemble methods, and evolutionary computation) may be the approaches that will bring the most technological progress in the coming years.

The divergence among the different factions within the AI community stems from the fact that no AI approach has been able to “solve” AI and dominate all other approaches. Each new approach we covered emerged from trying to address weaknesses in the then-current approaches. Cybernetics and control theory (Approach 1), for example, could only deal with problems that were encoded as differential equations. In response, heuristic problem-solving (Approach 2) aimed to encompass a broader set of problems—those that could be represented in terms of search on a state-space. In turn, expert systems and knowledge representation (Approach 5) aimed to extend the previous approaches by representing a larger set of problems—those that can be described using predicate logic. Connectionism and machine learning (Approach 6) allowed for dealing with problems that could not be described with predicate logic but could be represented in statistical terms and so on. Appreciating the strengths and weaknesses of different approaches highlights the vitality and dynamic nature of the AI field and suggests that, when a new AI approach emerges, it makes sense for OT to revisit old AI–OT analogies and to consider whether new analogies have become possible.

How AI May Continue to Inform OT Research

An important historical insight about the evolution of AI is Moravec’s (1988: 18) paradox:

It is comparatively easy to make computers exhibit adult-level performance in solving problems on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility.

In other words, what one would think is hard to program is easy and vice versa. Moravec’s observation was true when he made it, in 1988, because, at that point, most AI accomplishments had been

achieved using search approaches (Approaches 1–4), to which problems of perception and mobility had proven impervious. Only during the last decade have connectionist approaches (Approach 6) finally allowed computers to perform well in perception and mobility tasks—to graduate from world chess champion to average toddler.

Moravec’s (1988) observation is relevant to OT because the bulk of OT’s borrowings from AI relate to early search approaches. In fact, many ideas in the behavioral theory of the firm—such as aspiration levels, satisficing, and problemistic search—stem from search approaches, as do many ensuing ideas such as routines and search on rugged landscapes. But, in OT, we haven’t made an effort of a magnitude similar to the one made by the behavioral theory of the firm to integrate the AI ideas that came after the search approaches. For example, connectionism, reinforcement learning, and ensembles—all of which produced massive literatures and substantial progress in AI—do not have similarly massive counterparts in the OT literature. Of course, we have done some borrowing from those approaches, described above under their respective subsections, but there is probably more to be borrowed in light of the range and import of problems these approaches have been able to address in AI.

In the spirit of exploring how AI ideas may continue to inform OT, we suggest six possibilities, including two from each of the three families of approaches: search, representation, and aggregation.

Search: Understanding evolutionary economics in terms of genetic programming. “Evolutionary economics” (Nelson & Winter, 1982) studies firm adaptation, using economic models in which agents cannot optimize but instead adjust their decisions based on feedback. Such models are a cross between microeconomics and control theory and sometimes take the form of detailed simulations of specific cases, called “history-friendly models” (Malerba, Nelson, Orsenigo, & Winter, 2016). Although these models have provided many insights, evolutionary computation, which got off the ground a few years after Nelson and Winter’s book (Holland, 1992: ix), may provides a new set of analogies and tools with which to think more directly about evolutionary economics. In particular, modeling evolutionary economics using genetic programming (Koza, 1992) would more directly map key evolutionary processes in organizations, such as random variation, idea recombination, and the development of new technologies.

Search: Understanding organizational search in terms of the “subsumption architecture.” For the most part, search models in OT (e.g., Levinthal, 1997) have assumed that there is a unitary actor controlling the search and that the search follows a simple process, typically allowing for local search (hill-climbing) and some form of more distant search (such as imitation or random jumps). But most organizations have multiple conflicting goals, which may push the search in ways that are different from local and distant search.

An interesting development in AI that has not received much attention in OT is Brooks’s (1986) “subsumption architecture,” which proposed an alternative to the standard view of search in AI. Brooks’s idea was that a robot’s search in a landscape could be controlled by a hierarchy of processes, each one fixed on one objective—such as avoiding objects, wandering around, and exploring the world—in much the same way that firms pursue multiple and sometimes conflicting goals. Brooks’s idea could be used to model how organizations deal with conflict that emerges from receiving ambiguous performance feedback on multiple goals, which is a central yet understudied problem in OT (Gaba & Greve, 2019; Hu & Bettis, 2018).

Representation: Conceptualizing model-based search. Most of the literature on search in OT has been “representation free”; that is, the agent does not have a mental representation of the space being searched, but simply searches in the environment (for an exception, see Csaszar & Levinthal, 2016). (Interestingly, in AI, this type of search only became mainstream with Brooks’s [1991] “intelligence without representation.”) Yet it is a basic premise of OT that the environment is too vast and complex to be directly perceived by individuals, who therefore search a representation or simplified model of the environment and not the environment directly.

Two ideas from AI about searching with representations may therefore be fruitful in OT. One idea is search algorithms like A^* , which guide search not only by using the landscape’s fitness (which may be too expensive to consult often), but also by using a proxy of fitness that is more noisy but more available. The second idea is manifested in the programs that currently dominate backgammon, Go, and *StarCraft*, which combine representation with search. In these programs, representation is provided by a neural net that returns an approximate evaluation of any board configuration, whereupon search chooses which configurations to compare. This resembles the way individuals have an intuitive reaction that

can then be used to guide more deliberate search processes (i.e., System 1 and System 2; Evans, 2008). AI ideas such as these could add behavioral realism to the models of search used in OT.

Representation: Revisiting managerial mental maps. The edited volume by Huff (1990) described several ways of mapping the mental models of managers. Since then, two important representations have emerged in AI: Bayesian networks and neural nets. The success of these representations in AI attests to their ability to capture important aspects of the environment. It would seem natural, then, to try to map managers’ mental models using these newer forms of representation. Bayesian networks, for example, would allow capturing managers’ mental models about causality, an idea proposed by Durand and Vaara (2009) and Ryall (2009), but never, to the best of our knowledge, empirically validated in the strategy and organizations literatures. (Lee & Wagenmakers, 2013, illustrated this empirical method in the context of cognitive science.) Neural nets also provide a way of mapping managers’ cognition. For example, it would be possible to infer a manager’s “neural net” by having him or her play a business game; then, the neural net could be analyzed or used as an input to a simulation.

Aggregation: Using neural networks to understand bottom-up information processing. Most research on information aggregation has been on “flat” structures: small groups and crowds using rules like averaging or voting (e.g., Csaszar, 2013, 2019) and devoid of any structure. But most organizations have a hierarchical structure in which information is processed bottom up; information that enters through the lower levels flows up and may eventually reach the CEO.

This looks much like how information entering through the lowest layer of a neural net is processed and relayed while moving up through the net. It follows that neural nets might serve as an analogy of information aggregation in organizations. For example, one could use neural nets to study the effects of organizational characteristics such as hierarchy and span of control. Radner (1993) is among the few to have studied bottom-up information aggregation (see Garicano & Van Zandt, 2013), but his model made strong assumptions about the type of information processing the organization can perform. Neural nets provide a more general model of information processing in organizations.

Aggregation: Using ideas from distributed AI to understand information aggregation. When the OT literature on information aggregation (see, e.g.,

Csaszar & Eggers, 2013) has borrowed from AI, it has done so exclusively from the ensemble approach (Approach 10). This borrowing has allowed the OT literature on aggregation to analyze prediction problems—such as combining the opinions of multiple managers to estimate the quality of a project (Csaszar & Eggers, 2013)—which can be solved with ensembles. However, aggregation processes in organizations can be more involved than making a prediction. Imagine, for example, the aggregation process in group ideation tasks such as brainstorming a firm's strategy or writing a business plan. While ensembles are ill suited to model such tasks, distributed AI—by explicitly accounting for processes such as coordination, cooperation, and negotiation—seems to have the right attributes to model more complex aggregation tasks (for initial work along these lines, see Steinberger & Jung, 2019). Examining aggregation from the viewpoint of distributed AI may therefore be a fruitful approach for OT.

How AI Technologies May Affect Organizations

A fundamental idea in OT is bounded rationality: the premise that individual actors have hard limits to their ability to process information. AI alters that assumption by increasing those limits and thus expanding the bounds of rationality. This change in a key premise of OT cannot help but have many effects on organizations and much has already been written about such possible effects (Brynjolfsson & McAfee, 2014; Kretschmer & Khashabi, 2020; Raj & Seamans, 2019; von Krogh, 2018) and the ethical problems they may produce (Abbas, 2020; Möhlmann, Zalmanson, Henfridsson, & Gregory, 2020; Wallach & Allen, 2009). It is, however, out of the scope of this paper to review those works. Instead, we point out some possible scenarios and describe recent AI ideas on explainability and goal questioning (which someday may become AI approaches like the 10 we described) that may critically determine how AI will affect organizations.

Given the enormous progress that AI has made since its inception fewer than 70 years ago, the exponential increases in computing capacity, the increasing availability of data that can be used to train AI algorithms, and the increasing number of practitioners knowledgeable about AI, it seems more than likely that AI itself will continue to progress and that it will continue to be part of organizations and to change them. The future will therefore be somewhere in between these two scenarios, one in which

organizations use slightly more AI than today and one in which organizations are fully run by AIs.

We cannot predict what particular AI technologies will be used in this future, but we can predict that there will be greater demand for managers who know more about AI and who can design organizations and strategies that take advantage of it. We can also predict that there will be an increasing need to avoid the negative effects of AI; both regulators and ethicists should play an important role. For example, there could be a push toward AI technologies that can explain their decisions, rather than AIs that can just make good predictions (see, e.g., Hagrais, 2018). There could also be a push toward AI methods that can ensure their decisions comply with fairness criteria (for a survey on this topic, see Mehrabi, Morstatter, Saxena, Lerman, & Galstyan, 2019). Avoiding the negative effects of AI may also call for a different type of AI, one designed from the start to work for humans, rather than something designed to perform a task and then modified to avoid whatever human harm it turns out to cause. We may, in fact, need something akin to Asimov's (1942) three laws of robotics.

The AI technology to support this type of behavior does not yet exist, but one area of research that aims to address it is the attempt to produce AIs that are not certain about the objective function they are maximizing (Russell, 2019) and must therefore periodically seek confirmation from humans. This built-in uncertainty about their objective function would avoid the "paperclip maximizer" scenario (Bostrom, 2014: 150) in which an AI with the task of maximizing the production of paperclips would eventually eliminate humans, as they could consume resources useful for paperclip production and might even decide to turn off the AI and stop the all-important production of paperclips. With so much at stake, even if the chances of such a scenario are low, it is important that research examine this type of dubitative AI.

Closing

The aim of our paper has been to take stock of the many AI analogies used in OT and to serve as a launchpad for future exploration. Without closer attention to the AI–OT linkages, we in OT reduce our ability to draw on the AI field as a fertile source of ideas. And, as AI technologies become increasingly important for organizations, obscuring these linkages risks distancing OT theories from important areas of practice.

We hope that our paper has contributed to broadening the array of ideas that organization theorists can draw on when thinking about organizational processes. Our account of AI–OT linkages shows that they are much broader, deeper, and older than may be apparent in whatever AI technology is currently in the spotlight.

In terms of practical implications for OT researchers, we believe that they would benefit from continuing to use and develop AI analogies. AI is much more than a technology for OT; it is a set of models about how organizations work. PhD students in OT would be well advised to take at least one AI course. At a minimum, the knowledge gained in a machine-learning course will be useful as a research tool and will provide an understanding of the connectionist approaches. Better yet would be to acquire a broad overview of AI, which will help with theorizing—both to generate ideas and to describe them precisely.

The quest for AI—how to create a machine that thinks—is one of the great intellectual odysseys of our time. It has involved many fine scientists and produced much in terms of ideas and technologies. Staying connected to this quest is energizing. It is also an almost-bottomless well of ideas about how organizations can produce intelligent behavior.

REFERENCES

- Abbas, A. E. (Ed.). 2020. *Next-generation ethics: Engineering a better society*. Cambridge, U.K.: Cambridge University Press.
- Adami, C. 1998. *Introduction to artificial life*. New York, NY: Springer Science & Business Media.
- Agrawal, A., Gans, J., & Goldfarb, A. 2018. *Prediction machines: The simple economics of artificial intelligence*. Cambridge, MA: Harvard Business Press.
- Allen, P., Maguire, S., & McKelvey, B. (Eds.). 2011. *The SAGE handbook of complexity and management*. London, U.K.: SAGE.
- Anderson, P. 1999. Perspective: Complexity theory and organization science. *Organization Science*, 10: 216–232.
- Ansoff, H. I. 1965. *Corporate strategy*. New York, NY: McGraw-Hill.
- Aristotle. 1984. Prior analytics. In J. Barnes (Ed.), *Complete works of Aristotle: The revised Oxford translation*, vol. 1: 39–113. Princeton, NJ: Princeton University Press. (Original work composed c.350 BCE)
- Arthur, W. B. 1994. Inductive reasoning and bounded rationality. *American Economic Review*, 84: 406–411.
- Ashby, W. R. 1952. *Design for a brain*. New York, NY: Wiley.
- Ashby, W. R. 1956. *Introduction to cybernetics*. New York, NY: Wiley.
- Asimov, I. 1942. Runaround. *Astounding Science Fiction*, XXIX: 94–103.
- Axelrod, R. M. 1984. *The evolution of cooperation*. New York, NY: Basic Books.
- Axelrod, R. M., & Cohen, M. D. 2001. *Harnessing complexity*. New York, NY: Basic Books.
- Bailey, D. E., Faraj, S., Hinds, P., von Krogh, G., & Leonard, P. M. 2019. Emerging technologies and organizing. *Organization Science*, 30: 642–646.
- Bak, P., Tang, C., & Wiesenfeld, K. 1987. Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters*, 59: 381–384.
- Baldwin, C. Y. 2018. Bottlenecks, modules, and dynamic architectural capabilities. In D. J. Teece & S. Heaton (Eds.), *The Oxford handbook of dynamic capabilities*. Oxford, U.K.: Oxford University Press. doi: 10.1093/oxfordhb/9780199678914.013.011
- Baldwin, C. Y., & Clark, K. B. 2000. *The power of modularity*, vol. 1—design rules. Cambridge, MA: MIT Press.
- Barley, S. R. 1986. Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Administrative Science Quarterly*, 31: 78–108.
- Bateson, G. 1972. *Steps to an ecology of mind: Collected essays in anthropology, psychiatry, evolution, and epistemology*. San Francisco, CA: Chandler Publishing.
- Baum, J., & Haveman, H. A. 2020. Editors' comments: The future of organizational theory. *Academy of Management Review*, 45: 268–272.
- Baumann, O., Schmidt, J., & Stieglitz, N. 2019. Effective search on rugged performance landscapes: A review and outlook. *Journal of Management*, 45: 285–318.
- Bavelas, A. 1950. Communication patterns in task-oriented groups. *Journal of the Acoustical Society of America*, 22: 723–730.
- Bechky, B. A. 2003. Sharing meaning across occupational communities: The transformation of understanding on a production floor. *Organization Science*, 14: 312–330.
- Beer, S. 1972. *Brain of the firm: The managerial cybernetics of organization*. New York, NY: John Wiley & Sons.
- Bellman, R. 1957. *Dynamic programming*. Princeton, NJ: Princeton University Press.
- Bernoulli, J. 2006. *The art of conjecturing*. Baltimore, MD: Johns Hopkins University Press. (Original work published 1713)

- Bettis, R. A., & Blettner, D. P. 2020. Strategic reality today: Extraordinary past success, but difficult challenges loom. *Strategic Management Review*, 1: 75–101.
- Bingham, C. B., & Eisenhardt, K. M. 2011. Rational heuristics: The “simple rules” that strategists learn from process experience. *Strategic Management Journal*, 32: 1437–1464.
- Bobrow, D. G., & Norman, D. A. 1975. Some principles of memory schemata. In D. G. Bobrow & A. M. Collins (Eds.), *Representation and understanding: Studies in cognitive science*: 131–149. New York, NY: Academic Press.
- Bond, A. H., & Gasser, L. G. 1988. *Readings in distributed artificial intelligence*. San Francisco, CA: Morgan Kaufmann.
- Boole, G. 1951. *An investigation of the laws of thought*. New York, NY: Dover Publications. (Original work published 1854)
- Bostrom, N. 2014. *Superintelligence: Paths, dangers, strategies*. Oxford, U.K.: Oxford University Press.
- Braitenberg, V. 1984. *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.
- Breiman, L. 1996. Bagging predictors. *Machine Learning*, 24: 123–140.
- Brooks, R. 1986. A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, 2: 14–23.
- Brooks, R. A. 1991. Intelligence without representation. *Artificial Intelligence*, 47: 139–159.
- Brown, J. S., & Duguid, P. 1991. Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science*, 2: 40–57.
- Brown, S. L., & Eisenhardt, K. M. 1998. *Competing on the edge: Strategy as structured chaos*. Boston, MA: Harvard Business School Press.
- Bruderer, E., & Singh, J. V. 1996. Organizational evolution, learning, and selection: A genetic-algorithm-based model. *Academy of Management Journal*, 39: 1322–1349.
- Brusoni, S., Prencipe, A., & Pavitt, K. 2001. Knowledge specialization, organizational coupling, and the boundaries of the firm: Why do firms know more than they make? *Administrative Science Quarterly*, 46: 597–621.
- Brynjolfsson, E., & McAfee, A. 2014. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. New York, NY: W. W. Norton & Company.
- Buchanan, B., Sutherland, G., & Feigenbaum, E. A. 1969. Heuristic DENDRAL: A program for generating explanatory hypotheses in organic chemistry. In B. Meltzer & D. Michie (Eds.), *Machine intelligence*, vol. 4: 209–254. New York, NY: Edinburgh University Press.
- Burks, A. W. (Ed.). 1970. *Essays on cellular automata*. Urbana, IL: University of Illinois Press.
- Burt, R. S. 1992. *Structural holes: The social structure of competition*. Cambridge, MA: Harvard University Press.
- Burton, R. M., & Forsyth, J. D. 1986. Variety and the firm's performance: An empirical investigation. *Technovation*, 5: 9–21.
- Burton, R. M., & Obel, B. 2004. *Strategic organizational diagnosis and design: The dynamics of fit* (3rd ed.). Boston, MA: Kluwer.
- Bush, R. R., & Mosteller, F. 1955. *Stochastic models for learning*. New York, NY: Wiley.
- Campbell, M., Hoane, J., & Hsu, F. 2002. Deep Blue. *Artificial Intelligence*, 134: 57–83.
- Carley, K. M., & Gasser, L. 1999. Computational organization theory. In G. Weiss (Ed.), *Multiagent systems: A modern approach to distributed artificial intelligence*: 299–330. Cambridge, MA: MIT Press.
- Carley, K., & Newell, A. 1994. The nature of the social agent. *Journal of Mathematical Sociology*, 19: 221–262.
- Carley, K. M., & Prietula, M. J. 1994. *Computational Organization Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Carlile, P. R. 2002. A pragmatic view of knowledge and boundaries: Boundary objects in new product development. *Organization Science*, 13: 442–455.
- Carroll, T., & Burton, R. M. 2000. Organizations and complexity: Searching for the edge of chaos. *Computational & Mathematical Organization Theory*, 6: 319–337.
- Charniak, E., & McDermott, D. V. 1985. *Introduction to artificial intelligence*. Reading, MA: Addison-Wesley.
- Chinchalkar, S. 1996. An upper bound for the number of reachable positions. *ICGA Journal*, 19: 181–183.
- Choudhury, P., Allen, R. T., & Endres, M. G. 2021. Machine learning for pattern discovery in management research. *Strategic Management Journal*, 42: 30–57.
- Christensen, M., & Knudsen, T. 2010. Design of decision-making organizations. *Management Science*, 56: 71–89.
- Clarkson, G. P. E., & Simon, H. A. 1960. Simulation of individual and group behavior. *American Economic Review*, 50: 920–932.
- Cohen, M. D. 1981. The power of parallel thinking. *Journal of Economic Behavior & Organization*, 2: 285–306.

- Cohen, M. D., & Bacdayan, P. 1994. Organizational routines are stored as procedural memory: Evidence from a laboratory study. *Organization Science*, 5: 554–568.
- Condorcet, J. A. N. 1994. *Foundations of social choice and political theory* (I. McLean & F. Hewitt, Trans., Eds.). Brookfield, VT: Edgar Elgar. (Original work published 1785)
- Cooper, G. F., & Herskovits, E. 1992. A Bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9: 309–347.
- Crick, F. 1989. The recent excitement about neural networks. *Nature*, 337: 129–132.
- Crowston, K. 1992. Modeling coordination in organizations. In M. Masuch & M. Warglien (Eds.), *Artificial intelligence in organization and management theory*: 215–234. Amsterdam, Netherlands: Elsevier Science.
- Crutchfield, J. P., & Mitchell, M. 1995. The evolution of emergent computation. *Proceedings of the National Academy of Sciences of the United States of America*, 92: 10742–10746.
- Csaszar, F. A. 2012. Organizational structure as a determinant of performance: Evidence from mutual funds. *Strategic Management Journal*, 33: 611–632.
- Csaszar, F. A. 2013. An efficient frontier in organization design: Organizational structure as a determinant of exploration and exploitation. *Organization Science*, 24: 1083–1101.
- Csaszar, F. A. 2018. What makes a decision strategic? Strategic representations. *Strategy Science*, 3: 606–619.
- Csaszar, F. A. 2019. Limits to the wisdom of the crowd in idea selection. *Advances in Strategic Management*, 40: 275–297.
- Csaszar, F. A., & Eggers, J. P. 2013. Organizational decision making: An information aggregation view. *Management Science*, 59: 2257–2277.
- Csaszar, F. A., & Levinthal, D. A. 2016. Mental representation and the discovery of new strategies. *Strategic Management Journal*, 37: 2031–2049.
- Csaszar, F. A., & Ostler, J. 2020. A contingency theory of representational complexity in organizations. *Organization Science*, 31: 1198–1219.
- Cybenko, G. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2: 303–314.
- Cyert, R. M., Feigenbaum, E. A., & March, J. G. 1959. Models in a behavioral theory of the firm. *Behavioral Science*, 4: 81–95.
- Cyert, R. M., & March, J. G. 1955. Organizational structure and pricing behavior in an oligopolistic market. *American Economic Review*, 45: 129–139.
- Cyert, R. M., & March, J. G. 1963. *A behavioral theory of the firm*. Englewood Cliffs, NJ: Prentice Hall.
- Dasgupta, S. 2003. Multidisciplinary creativity: The case of Herbert A. Simon. *Cognitive Science*, 27: 683–707.
- de Finetti, B. 1980. Foresight: Its logical laws, its subjective sources. In H. E. Kyburg & H. E. Smokler (Eds.), *Studies in subjective probability*: 53–118. Huntington, NY: Robert E. Krieger. (Original work published 1937)
- Dennett, D. C. 1984. Cognitive wheels: The frame problem of AI. In C. Hookway (Ed.), *Minds, machines and evolution: Philosophical studies*: 129–151. Cambridge, U.K.: Cambridge University Press.
- Denrell, J., Fang, C., & Levinthal, D. A. 2004. From T-mazes to labyrinths: Learning from model-based feedback. *Management Science*, 50: 1366–1378.
- Denrell, J., & March, J. G. 2001. Adaptation as information restriction: The hot stove effect. *Organization Science*, 12: 523–538.
- Descartes, R. 1985. Treatise on man. *The philosophical writings of Descartes* (J. Cottingham, R. Stoothoff, & D. Murdoch, Trans.), vol. I: 99–108. Cambridge, U.K.: Cambridge University Press. (Original work published 1633)
- Dietterich, T. G. 2000. Ensemble methods in machine learning. In G. Kittler & F. Roli (Eds.), *Multiple classifier systems*, vol. 1857: 1–15. Berlin, Germany: Springer.
- DiMaggio, P. J., & Powell, W. W. 1983. The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48: 147–160.
- Dreyfus, H. L. 1972. *What computers can't do: A critique of artificial reason*. New York, NY: Harper & Row.
- Dreyfus, H. L., & Dreyfus, S. E. 1986. *Mind over machine*. New York, NY: Free Press.
- Dreyfus, H. L., & Dreyfus, S. E. 1988. Making a mind versus modeling the brain: Artificial intelligence back at a branchpoint. In S. Graubard (Ed.), *The artificial debate*: 15–43. Cambridge, MA: MIT Press.
- Duda, R. O., Hart, P. E., Nilsson, N. J., Reboh, R., Slocum, J., & Sutherland, G. L. 1977. *Development of a computer-based consultant for mineral exploration* (SRI International, annual report, SRI projects 5821 and 6415). Menlo Park, CA: Stanford Research Institute.
- Durand, R., & Vaara, E. 2009. Causation, counterfactuals, and competitive advantage. *Strategic Management Journal*, 30: 1245–1264.
- Edwards, D. J., & Hart, T. P. 1963. *The alpha-beta heuristic* (MIT AI memo AIM-030). Cambridge, MA: MIT Computer Science and Artificial Intelligence Laboratory.

- Efron, B., & Tibshirani, R. J. 1993. *An introduction to the bootstrap*. New York, NY: Chapman & Hall.
- Eisenhardt, K. M., & Piezunka, H. 2011. Complexity theory and corporate strategy. In P. Allen, S. Maguire, & B. McKelvey (Eds.), *The SAGE handbook of complexity and management*: 506–523. London, U.K.: SAGE.
- Ericsson, K. A., & Simon, H. A. 1992. *Protocol analysis: Verbal reports as data* (rev. ed.). Cambridge, MA: MIT Press.
- Erman, L. D., Hayes-Roth, F., Lesser, V. R., & Reddy, D. R. 1980. The Hearsay-II speech-understanding system: Integrating knowledge to resolve uncertainty. *ACM Computing Surveys*, 12: 213–253.
- Ethiraj, S. K., & Levinthal, D. A. 2004. Modularity and innovation in complex systems. *Management Science*, 50: 159–173.
- Evans, J. S. B. T. 2008. Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59: 255–278.
- Fang, C., Lee, J., & Schilling, M. A. 2010. Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science*, 21: 625–642.
- Fang, C., & Levinthal, D. A. 2009. Near-term liability of exploitation: Exploration and exploitation in multi-stage problems. *Organization Science*, 20: 538–551.
- Feigenbaum, E. A. 1992. A personal view of expert systems: Looking back and looking ahead. *Expert Systems with Applications*, 5: 193–201.
- Feigenbaum, E. A., & McCorduck, P. 1983. *Fifth generation: Artificial intelligence and Japan's computer challenge to the world*. Reading, MA: Addison-Wesley.
- Feigenbaum, E. A., McCorduck, P., & Nii, H. P. 1988. *The rise of the expert company*. New York, NY: Times Books.
- Feldman, M. S., & Pentland, B. T. 2003. Reconceptualizing organizational routines as a source of flexibility and change. *Administrative Science Quarterly*, 48: 94–118.
- Felin, T., Kauffman, S., Koppl, R., & Longo, G. 2014. Economic opportunity and evolution: Beyond landscapes and bounded rationality. *Strategic Entrepreneurship Journal*, 8: 269–282.
- Flake, G. W. 1998. *The computational beauty of nature: Computer explorations of fractals, chaos, complex systems, and adaptation*. Cambridge, MA: MIT Press.
- Floreano, D., & Mattiussi, C. 2008. *Bio-inspired artificial intelligence: Theories, methods, and technologies*. Cambridge, MA: MIT Press.
- Flores, F., Graves, M., Hartfield, B., & Winograd, T. 1988. Computer systems and the design of organizational interaction. *ACM Transactions on Information Systems*, 6: 153–172.
- Fogel, D. B. (Ed.). 1998. *Evolutionary computation: The fossil record*. Piscataway, NJ: IEEE Press.
- Forrester, J. W. 1961. *Industrial dynamics*. Cambridge, MA: MIT Press.
- Freeman, J., Larsen, E. R., & Lomi, A. 2012. Why is there no cannery in “Cannery Row”? Exploring a behavioral simulation model of population extinction. *Industrial and Corporate Change*, 21: 99–125.
- Frege, G. 1967. Begriffsschrift: A formal language of pure thought modelled on that of arithmetic. In J. van Heijenoort (Ed.), *From Frege to Gödel: A source book in mathematical logic, 1879–1931*: 1–82. Cambridge, MA: Harvard University Press. (Original work published 1879)
- Freud, S. 1933. *New introductory lectures on psychoanalysis*. New York, NY: W. W. Norton & Company.
- Freund, Y., & Schapire, R. 1996. Experiments with a new boosting algorithm. In L. Saitta (Ed.), *Proceedings of the thirteenth international conference on machine learning*: 148–156. San Francisco, CA: Morgan Kaufmann.
- Friedman, N., Linial, M., Nachman, I., & Pe'er, D. 2000. Using Bayesian networks to analyze expression data. *Journal of Computational Biology*, 7: 601–620.
- Gaba, V., & Greve, H. R. 2019. Safe or profitable? The pursuit of conflicting goals. *Organization Science*, 30: 647–667.
- Galbraith, J. R. 1973. *Designing complex organizations*. Reading, MA: Addison-Wesley.
- Galton, F. 1949. Vox populi. *Nature*, 75: 450–451. (Original work published 1907)
- Gardner, M. 1970. Mathematical games. *Scientific American*, 222: 132–140.
- Garicano, L., & Van Zandt, T. 2013. Hierarchies and the division of labor. In R. Gibbons & J. Roberts (Eds.), *The handbook of organizational economics*: 604–654. Princeton, NJ: Princeton University Press.
- Gary, M. S., Kunc, M., Morecroft, J. D., & Rockart, S. F. 2008. System dynamics and strategy. *System Dynamics Review*, 24: 407–429.
- Gavetti, G., & Warglien, M. 2015. A model of collective interpretation. *Organization Science*, 26: 1263–1283.
- Geman, S., Bienenstock, E., & Doursat, R. 1992. Neural networks and the bias/variance dilemma. *Neural Computation*, 4: 1–58.
- Ginsberg, M. (Ed.). 1987. *Readings in nonmonotonic reasoning*. San Francisco, CA: Morgan Kaufmann.
- Gioia, D. A., & Poole, P. P. 1984. Scripts in organizational behavior. *Academy of Management Review*, 9: 449–459.

- Graefe, A., Küchenhoff, H., Stierle, V., & Riedl, B. 2015. Limitations of ensemble Bayesian model averaging for forecasting social science problems. *International Journal of Forecasting*, 31: 943–951.
- Grant, R. M. 1996. Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17: 109–122.
- Greve, H. R. 1998. Performance, aspirations, and risky organizational change. *Administrative Science Quarterly*, 43: 58–86.
- Grushka-Cockayne, Y., Jose, V. R. R., & Lichtendahl, K. C. 2017. Ensembles of overfit and overconfident forecasts. *Management Science*, 63: 1110–1130.
- Hacking, I. 2006. *The emergence of probability: A philosophical study of early ideas about probability, induction and statistical inference* (2nd ed.). Cambridge, U.K.: Cambridge University Press.
- Hagras, H. 2018. Toward human-understandable, explainable AI. *Computer*, 51: 28–36.
- Hambrick, D. C., & Mason, P. A. 1984. Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review*, 9: 193–206.
- Hannan, M. T., & Freeman, J. 1989. *Organizational ecology*. Cambridge, MA: Ballinger.
- Hannan, M. T., Polos, L., & Carroll, G. R. 2007. *Logics of organization theory: Audiences, codes, and ecologies*. Princeton, NJ: Princeton University Press.
- Hansen, L. K., & Salamon, P. 1990. Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12: 993–1001.
- Hanson, J. E. 2009. Emergent phenomena in cellular automata. In R. A. Meyers (Ed.), *Encyclopedia of complexity and systems science*: 768–778. New York, NY: Springer.
- Hargadon, A. B., & Bechky, B. A. 2006. When collections of creatives become creative collectives: A field study of problem solving at work. *Organization Science*, 17: 484–500.
- Hart, P. E., Nilsson, N. J., & Raphael, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4: 100–107.
- Hasson, U., Nastase, S. A., & Goldstein, A. 2020. Direct fit to nature: An evolutionary perspective on biological and artificial neural networks. *Neuron*, 105: 416–434.
- Hayes-Roth, F. 1992. Expert systems. In S. C. Shapiro (Ed.), *Encyclopedia of artificial intelligence*: 287–298. New York, NY: Wiley & Sons.
- Hebb, D. O. 1949. *The organization of behavior: A neuropsychological theory*. New York, NY: Wiley.
- Hegselmann, R. 1996. Cellular automata in the social sciences. In R. Hegselmann, U. Mueller, & K. G. Troitzsch (Eds.), *Modelling and simulation in the social sciences from the philosophy of science point of view*: 209–233. Dordrecht, Netherlands: Springer.
- Hernes, T., & Bakken, T. 2003. Implications of self-reference: Niklas Luhmann's autopoiesis and organization theory. *Organization Studies*, 24: 1511–1535.
- Hewitt, C. 1977. Viewing control structures as patterns of passing messages. *Artificial Intelligence*, 8: 323–364.
- Hewitt, C. 1988. Offices are open systems. *ACM Transactions on Office Information Systems*, 4: 271–287.
- Ho, T. K. 1995. Random decision forests. In Institute of Electrical and Electronics Engineers (Ed.), *Proceedings of the third international conference on document analysis and recognition*: 278–282. Los Alamitos, CA: IEEE Computer Society.
- Holland, J. H. 1975. *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. Ann Arbor, MI: University of Michigan Press.
- Holland, J. H. 1992. *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence* (2nd ed.). Cambridge, MA: MIT Press.
- Holland, J. H., & Miller, J. H. 1991. Artificial adaptive agents in economic theory. *American Economic Review*, 81: 365–370.
- Hopfield, J. J. 1982. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, 79: 2554–2558.
- Hornik, K., Stinchcombe, M., & White, H. 1989. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2: 359–366.
- Hu, S., & Bettis, R. A. 2018. Multiple organization goals with feedback from shared technological task environments. *Organization Science*, 29: 873–889.
- Huberman, B. A. 1988. *Ecology of computation*. Amsterdam, Netherlands: North-Holland.
- Huff, A. S. 1990. *Mapping strategic thought*. Chichester, U.K.: Wiley.
- Jordan, M. I., & Mitchell, T. M. 2015. Machine learning: Trends, perspectives, and prospects. *Science*, 349: 255–260.
- Joseph, J., & Gaba, V. 2020. Organizational structure, information processing, and decision-making: A retrospective and road map for research. *Academy of Management Annals*, 14: 267–302.
- Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., Osawa, E., & Matsubara, H. 1997. RoboCup: A challenge problem for AI. *AI Magazine*, 18: 73–85.

- Knudsen, T., & Levinthal, D. A. 2007. Two faces of search: Alternative generation and alternative evaluation. *Organization Science*, 18: 39–54.
- Kogut, B., & Zander, U. 1992. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science*, 3: 383–397.
- Koller, D., & Friedman, N. 2009. *Probabilistic graphical models: Principles and techniques*. Cambridge, MA: MIT Press.
- Koren, Y. 2009. *The BellKor solution to the Netflix Grand Prize* (Netflix Prize documentation 81). Los Gatos, CA: Netflix.
- Koza, J. R. 1992. *Genetic programming: On the programming of computers by means of natural selection*. Cambridge, MA: MIT Press.
- Koza, J. R. 2010. Human-competitive results produced by genetic programming. *Genetic Programming and Evolvable Machines*, 11: 251–284.
- Kretschmer, T., & Khashabi, P. 2020. Digital transformation and organization design: An integrated approach. *California Management Review*, 62: 86–104.
- Kuncheva, L. I. 2014. *Combining pattern classifiers: Methods and algorithms* (2nd ed.). Hoboken, NJ: Wiley.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. 1987. SOAR: An architecture for general intelligence. *Artificial Intelligence*, 33: 1–64.
- Langton, C. 1986. Studying artificial life with cellular automata. *Physica D. Nonlinear Phenomena*, 22: 120–149.
- Laureiro-Martínez, D., Brusoni, S., Canessa, N., & Zollo, M. 2015. Understanding the exploration–exploitation dilemma: An fMRI study of attention control and decision-making performance. *Strategic Management Journal*, 36: 319–338.
- Lave, C. A., & March, J. G. 1975. *An introduction to models in the social sciences*. New York, NY: Harper & Row.
- Lee, J., Lee, K., & Rho, S. 2002. An evolutionary perspective on strategic group emergence: A genetic algorithm-based model. *Strategic Management Journal*, 23: 727–746.
- Lee, M. D., & Wagenmakers, E.-J. 2013. *Bayesian cognitive modeling: A practical course*. Cambridge, U.K.: Cambridge University Press.
- Lenat, D. B., Prakash, M., & Shepherd, M. 1985. CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine*, 6: 65–85.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science*, 43: 934–950.
- Liang, D. W., Moreland, R., & Argote, L. 1995. Group versus individual training and group-performance: The mediating role of transactive memory. *Personality and Social Psychology Bulletin*, 21: 384–393.
- Lindsay, R. K., Buchanan, B. G., Feigenbaum, E. A., & Lederberg, J. 1993. DENDRAL: A case study of the first expert system for scientific hypothesis formation. *Artificial Intelligence*, 61: 209–261.
- Lippman, S. A., & Rumelt, R. P. 1982. Uncertain imitability: An analysis of interfirm differences in efficiency under competition. *Bell Journal of Economics*, 13: 418–438.
- Lomi, A., & Larsen, E. R. 1996. Interacting locally and evolving globally: A computational approach to the dynamics of organizational populations. *Academy of Management Journal*, 39: 1287–1321.
- Lovullo, D., & Kahneman, D. 2003. Delusions of success: How optimism undermines executives' decisions. *Harvard Business Review*, 81: 56–63.
- Lovelace, A. 1843. Sketch of the analytical engine invented by Charles Babbage—notes upon the memoir by the translator. In R. Taylor (Ed.), *Scientific memoirs*, Vol. 3: 666–731. London, U.K.: Richard Taylor.
- Luger, G. F. 2005. *Artificial intelligence: Structures and strategies for complex problem solving* (5th ed.). Boston, MA: Pearson Education.
- Luger, G. F., & Stubblefield, W. A. 1993. *Artificial intelligence: Structures and strategies for complex problem solving* (2nd ed.). Redwood City, CA: Benjamin/Cummings.
- Luhmann, N. 1990. *Essays on self-reference*. New York, NY: Columbia University Press.
- Lustick, I. S. 2000. Agent-based modelling of collective identity: Testing constructivist theory. *Journal of Artificial Societies and Social Simulation*, 3. Retrieved from <https://www.jasss.org/3/1/1.html>
- Malerba, F., Nelson, R. R., Orsenigo, L., & Winter, S. G. 2016. *Innovation and the evolution of industries: History-friendly models*. Cambridge, U.K.: Cambridge University Press.
- Malone, T. W., & Crowston, K. 1991. *Toward an interdisciplinary theory of coordination* (MIT Center for Coordination Science, technical report 120). Cambridge, MA: Center for Coordination Science, Sloan School of Management, Massachusetts Institute of Technology.
- Malone, T. W., Crowston, K., & Herman, G. A. 2003. *Organizing business knowledge: The MIT process handbook*. Cambridge, MA: MIT Press.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science*, 2: 71–87.
- March, J. G. 1999. *The pursuit of organizational intelligence*. Malden, MA: Blackwell Publishers.
- March, J. G., & Simon, H. A. 1958/1993. *Organizations*. New York, NY: John Wiley & Co. (Original work published 1958)

- Marchiori, D., & Warglien, M. 2008. Predicting human interactive learning by regret-driven neural networks. *Science*, 319: 1111–1113.
- Maturana, H. R., & Varela, F. J. 1980. *Autopoiesis and cognition: The realization of the living*. Dordrecht, Netherlands: Reidel.
- Mayor, A. 2018. *Gods and robots: Myths, machines, and ancient dreams of technology*. Princeton, NJ: Princeton University Press.
- McCarthy, J. 1980. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13: 27–39.
- McCarthy, J., & Hayes, P. J. 1981. Some philosophical problems from the standpoint of artificial intelligence. In B. L. Webber & N. J. Nilsson (Eds.), *Readings in artificial intelligence*: 431–450. San Francisco, CA: Morgan Kaufmann. (Original work published 1969)
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. 2006. A proposal for the Dartmouth summer research project on artificial intelligence. *AI Magazine*, 27: 12–15. (Original work published 1955)
- McCorduck, P. 2004. *Machines who think: A personal inquiry into the history and prospects of artificial intelligence* (25th anniversary ed.). Natick, MA: A. K. Peters.
- McCulloch, W. S., & Pitts, W. 1943. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5: 115–133.
- Meadows, D. H., & Club of Rome. 1972. *Limits to growth: A report for the Club of Rome's project on the predicament of mankind*. New York, NY: Universe Books.
- Medina, E. 2011. *Cybernetic revolutionaries: Technology and politics in Allende's Chile*. Cambridge, MA: MIT Press.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. 2019. A survey on bias and fairness in machine learning. Retrieved from <https://arxiv.org/abs/1908.09635>
- Miller, J. H., & Page, S. E. 2007. *Complex adaptive systems: An introduction to computational models of social life*. Princeton, NJ: Princeton University Press.
- Mindell, D. 2002. *Between human and machine: Feedback, control, and computing before cybernetics*. Baltimore, MD: Johns Hopkins University Press.
- Minsky, M. L. 1954. *Theory of neural-analog reinforcement systems and its application to the brain-model problem* (PhD thesis). Princeton University, Princeton, NJ.
- Minsky, M. L. 1967. *Computation: Finite and infinite machines*. Englewood Cliffs, NJ: Prentice-Hall.
- Minsky, M. L. 1975. A framework for representing knowledge. In P. Winston (Ed.), *The psychology of computer vision*: 211–277. New York, NY: McGraw-Hill.
- Minsky, M. L. 1986. *The society of mind*. New York, NY: Simon and Schuster.
- Minsky, M. L., & Papert, S. A. 1969. *Perceptrons: An introduction to computational geometry*. Cambridge, MA: MIT Press.
- Mitchell, M. 1996. *An introduction to genetic algorithms*. Cambridge, MA: MIT Press.
- Mitchell, M. 1998. Computation in cellular automata: A selected review. In T. Gramß, S. Bornholdt, M. Groß, M. Mitchell, & T. Pellizzari (Eds.), *Non-standard computation*: 95–140. Weinheim, Germany: Wiley-VCH Verlag.
- Mitchell, M. 2009. *Complexity: A guided tour*. New York, NY: Oxford University Press.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. 2013. Playing Atari with deep reinforcement learning. Retrieved from <https://arxiv.org/abs/1312.5602>
- Möhlmann, M., Zalmanson, L., Henfridsson, O., & Gregory, R. W. 2020. Algorithmic management of work on online labor platforms: When matching meets control. *Management Information Systems Quarterly*. doi: 10.25300/MISQ/2021/15333
- Moore, E. F., & Shannon, C. E. 1993. Reliable circuits using less reliable relays (I and II). In N. J. A. Sloane & A. D. Wyner (Eds.), *Claude Elwood Shannon: Collected papers*: 796–830. New York, NY: Wiley-IEEE Press. (Original work published 1956)
- Moravec, H. P. 1988. *Mind children: The future of robot and human intelligence*. Cambridge, MA: Harvard University Press.
- Morgan, G. 2006. *Images of organization*. Thousand Oaks, CA: SAGE.
- Moses, J. 1974. MACSYMA—the fifth year. *ACM SIGSAM Bulletin*, 8: 105–110.
- Nelson, R. R., & Winter, S. G. 1982. *An evolutionary theory of economic change*. Cambridge, MA: Harvard University Press.
- Newell, A. 1983. The heuristic of George Polya and its relation to artificial intelligence. In R. Groner, M. Groner, & W. Bischof (Eds.), *Methods of heuristics*: 195–243. Hillsdale, NJ: Erlbaum.
- Newell, A., Shaw, J. C., & Simon, H. A. 1958. Chess-playing programs and the problem of complexity. *IBM Journal of Research and Development*, 2: 320–335.
- Newell, A., Shaw, J. C., & Simon, H. A. 1959. Report on a general problem-solving program. In United Nations Educational, Scientific and Cultural Organisation (Ed.), *Proceedings of the international conference on information processing*: 256–264. Paris, France: UNESCO.
- Newell, A., & Simon, H. A. 1956. The logic theory machine: A complex information processing system. *I.R.E. Transactions on Information Theory*, 2: 61–79.

- Newell, A., & Simon, H. A. 1972. *Human problem solving*. Englewood Cliffs, NJ: Prentice Hall.
- Newell, A., & Simon, H. A. 1976. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19: 113–126.
- Nickerson, J. A., & Zenger, T. R. 2002. Being efficiently fickle: A dynamic theory of organizational choice. *Organization Science*, 13: 547–566.
- Nilsson, N. J. 1965. *Learning machines: Foundations of trainable pattern-classifying systems*. New York, NY: McGraw-Hill.
- Nilsson, N. J. 1998. *Artificial intelligence: A new synthesis*. San Francisco, CA: Morgan Kaufmann.
- Nilsson, N. J. 2010. *The quest for artificial intelligence: A history of ideas and achievements*. Cambridge, U.K.: Cambridge University Press.
- Nonaka, I. 1994. A dynamic theory of organizational knowledge creation. *Organization Science*, 5: 14–37.
- Norvig, P. 1992. *Paradigms of artificial intelligence programming: Case studies in Common Lisp*. San Francisco, CA: Morgan Kaufmann.
- Nowak, A., & Lewenstein, M. 1996. Modeling social change with cellular automata. In R. Hegselmann, U. Mueller, & K. G. Troitzsch (Eds.), *Modelling and simulation in the social sciences from the philosophy of science point of view*: 249–285. Dordrecht, Netherlands: Springer.
- Nowak, M. A., & May, R. M. 1992. Evolutionary games and spatial chaos. *Nature*, 359: 826–829.
- Oaksford, M., & Chater, N. 2007. *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford, U.K.: Oxford University Press.
- Ocasio, W., Rhee, L., & Boynton, D. 2020. March and the pursuit of organizational intelligence: The interplay between procedural rationality and sensible foolishness. *Industrial and Corporate Change*, 29: 225–239.
- Orr, J. 1996. *Talking about machines: An ethnography of a modern job*. Ithaca, NY: ILR Press.
- Page, S. E. 2007. *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton, NJ: Princeton University Press.
- Page, S. E. 2018. *The model thinker: What you need to know to make data work for you*. New York, NY: Basic Books.
- Pavlov, P. I. 1927. *Conditioned reflexes*. London, U.K.: Oxford University Press.
- Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Francisco, CA: Morgan Kaufmann.
- Pearl, J. 2000. *Causality: Models, reasoning, and inference*. Cambridge, U.K.: Cambridge University Press.
- Pentland, B. T. 1995. Grammatical models of organizational processes. *Organization Science*, 6: 541–556.
- Pias, C. 2016. *Cybernetics: The Macy conferences 1946–1953—the complete transactions*. Zurich, Switzerland: Diaphanes.
- Pierce, W. 1961. *Improving reliability of digital systems by redundancy and adaptation* (Unpublished doctoral dissertation). Stanford University, Stanford, CA.
- Polos, L., Hannan, M. T., & Carroll, G. R. 2002. Foundations of a theory of social forms. *Industrial and Corporate Change*, 11: 85–115.
- Polya, G. 1945. *How to solve it: A new aspect of mathematical method*. Princeton, NJ: Princeton University Press.
- Pontikes, E. G. 2018. Category strategy for firm advantage. *Strategy Science*, 3: 620–631.
- Poole, D. L., & Mackworth, A. K. 2010. *Artificial intelligence: Foundations of computational agents*. New York, NY: Cambridge University Press.
- Posen, H., Keil, T., Kim, S., & Meissner, F. 2018. Renewing research on problemistic search: A review and research agenda. *Academy of Management Annals*, 12: 208–251.
- Posen, H. E., & Levinthal, D. A. 2012. Chasing a moving target: Exploitation and exploration in dynamic environments. *Management Science*, 58: 587–601.
- Prietula, M., Carley, K., & Gasser, L. (Eds.). 1998. *Simulating organizations: Computational models of institutions and groups*, vol. 1. Cambridge, MA: MIT Press.
- Prietula, M. J., & Simon, H. A. 1989. Experts in your midst. *Harvard Business Review*, 67: 120–124.
- Puranam, P., Stieglitz, N., Osman, M., & Pillutla, M. M. 2015. Modelling bounded rationality in organizations: Progress and prospects. *Academy of Management Annals*, 9: 337–392.
- Puranam, P., & Swamy, M. 2016. How initial representations shape coupled learning processes. *Organization Science*, 27: 323–335.
- Radner, R. 1993. The organization of decentralized information processing. *Econometrica*, 61: 1109–1146.
- Rahmandad, H. 2008. Effect of delays on complexity of organizational learning. *Management Science*, 54: 1297–1312.
- Raisch, S., & Birkinshaw, J. 2008. Organizational ambidexterity: Antecedents, outcomes, and moderators. *Journal of Management*, 34: 375–409.
- Raj, M., & Seamans, R. 2019. Primer on artificial intelligence and robotics. *Journal of Organization Design*, 8: 1–14.
- Ramsey, F. P. 1931. Truth and probability. In R. B. Braithwaite (Ed.), *The foundations of mathematics*

- and other logical essays:** 156–198. London, U.K.: Kegan Paul, Trench, Trubner, & Co.
- Ren, Y., & Argote, L. 2011. Transactive memory systems 1985–2010: An integrative framework of key dimensions, antecedents, and consequences. *Academy of Management Annals*, 5: 189–229.
- Repenning, N. P. 2002. A simulation-based approach to understanding the dynamics of innovation implementation. *Organization Science*, 13: 109–127.
- Rich, E. 1983. *Artificial intelligence*. New York, NY: McGraw-Hill.
- Rivkin, J. W. 2000. Imitation of complex strategies. *Management Science*, 46: 824–844.
- Rosenblatt, F. 1958. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65: 386–408.
- Rosenschein, J. S., & Genesereth, M. R. 1985. Deals among rational agents. In A. Joshi (Ed.), *Proceedings of the 9th international joint conference on artificial intelligence*, Vol. 1: 91–99. San Francisco, CA: Morgan Kaufmann.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986. Learning representations by back-propagating errors. *Nature*, 323: 533–536.
- Russell, S. J. 2019. *Human compatible: Artificial intelligence and the problem of control*. New York, NY: Viking.
- Russell, S. J., & Norvig, P. 2020. *Artificial intelligence: A modern approach* (4th ed.). Hoboken, NJ: Pearson.
- Ryall, M. D. 2009. Causal ambiguity, complexity, and capability-based advantage. *Management Science*, 55: 389–403.
- Ryall, M. D., & Bramson, A. L. 2014. *Inference and intervention: Causal models for business analysis*. New York, NY: Routledge.
- Sah, R. K., & Stiglitz, J. E. 1986. The architecture of economic systems: Hierarchies and polyarchies. *American Economic Review*, 76: 716–727.
- Samuel, A. L. 1959. Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3: 210–229.
- Sarasvathy, S. D. 2008. *Effectuation: Elements of entrepreneurial expertise*. Cheltenham, U.K.: Edward Elgar.
- Schank, R., & Abelson, R. P. 1977. *Scripts, plans, goals, and understanding*. Hillsdale, NJ: Erlbaum.
- Schapire, R. E. 1990. The strength of weak learnability. *Machine Learning*, 5: 197–227.
- Schelling, T. C. 1969. Models of segregation. *American Economic Review*, 59: 488–493.
- Schelling, T. C. 1971. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1: 143–186.
- Schiff, J. L. 2008. *Cellular automata: A discrete view of the world*. New York, NY: John Wiley & Sons.
- Schilling, M. A. 2000. Toward a general modular systems theory and its application to interfirm product modularity. *Academy of Management Review*, 25: 312–334.
- Scott, W. R., & Davis, G. F. 2007. *Organizations and organizing: Rational, natural, and open system perspectives*. Upper Saddle River, NJ: Pearson–Prentice Hall.
- Seidel, V. P., & O’Mahony, S. 2014. Managing the repertoire: stories, metaphors, prototypes, and concept coherence in product innovation. *Organization Science*, 25: 691–712.
- Seidl, D. 2004. *Luhmann’s theory of autopoietic social systems* (Munich Business Research working paper no. 2004-2). Munich, Germany: Munich School of Management, Ludwig Maximilian University of Munich.
- Sejnowski, T. J. 2018. *The deep learning revolution*. Cambridge, MA: MIT Press.
- Selfridge, O. 1959. Pandemonium: A paradigm for learning. In D. V. Blake & A. M. Uttley (Eds.), *Proceedings of the symposium on mechanisation of thought processes*: 511–531. London, U.K.: Her Majesty’s Stationery Office.
- Shannon, C. E. 1948. A mathematical theory of communication. *Bell System Technical Journal*, 27: 379–423.
- Shannon, C. E. 1950. Programming a computer for playing chess. *Philosophical Magazine*, 41: 256–275.
- Shannon, C. E., & Weaver, W. 1998. *The mathematical theory of communication*. Urbana, IL: University of Illinois Press. (Original work published 1949)
- Shoham, Y., & Leyton-Brown, K. 2009. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge, U.K.: Cambridge University Press.
- Shortliffe, E., Davis, R., Axline, S. G., Buchanan, B. G., Green, C. C., & Cohen, S. N. 1975. Computer-based consultations in clinical therapeutics: Explanation and rule acquisition capabilities of the MYCIN system. *Computers and Biomedical Research*, 8: 303–320.
- Siggelkow, N. 2002. Evolution toward fit. *Administrative Science Quarterly*, 47: 125–159.
- Siggelkow, N., & Levinthal, D. A. 2003. Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Organization Science*, 14: 650–669.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I.,

- Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529: 484–489.
- Simon, H. A. 1952. On the application of servomechanism theory in the study of production control. *Econometrica*, 20: 247–268.
- Simon, H. A. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69: 99–118.
- Simon, H. A. 1956. Rational choice and the structure of the environment. *Psychological Review*, 63: 129–138.
- Simon, H. A. 1965. *The shape of automation for men and management*. New York, NY: Harper & Row.
- Simon, H. A. 1996. *The sciences of the artificial* (3rd ed.). Cambridge, MA: MIT Press. (Original work published 1969)
- Simon, H. A., & Newell, A. 1958. Heuristic problem solving: The next advance in operations research. *Operations Research*, 6: 1–10.
- Skinner, B. F. 1938. *The behavior of organisms: An experimental analysis*. New York, NY: Appleton-Century.
- Spender, J. C. 2013. Herbert Alexander Simon: Philosopher of the organizational life-world. In M. Witzel & M. Warner (Eds.), *The Oxford handbook of management theorists*: 297–357. Oxford, U.K.: Oxford University Press.
- Spirtes, P., Glymour, C., & Scheines, R. 2000. *Causation, prediction, and search*. Cambridge, MA: MIT Press.
- Star, S. 1989. The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving. In L. Gasser & M. Huhns (Eds.), *Readings in distributed artificial intelligence*: 37–54. Menlo Park, CA: Morgan Kaufmann.
- Steinberger, T., & Jung, J. Y. 2019. Designing the microstructure of routines. *Journal of Organization Design*, 8: 1–18.
- Sterman, J. 2000. *Business dynamics: Systems thinking and modeling for a complex world*. Boston, MA: Irwin/McGraw-Hill.
- Suchman, L. A. 1987. *Plans and situated actions: The problem of human-machine communication*. Cambridge, U.K.: Cambridge University Press.
- Surowiecki, J. 2004. *The wisdom of crowds*. New York, NY: Doubleday.
- Sutton, R. 2019, March 13. The bitter lesson. Retrieved from <http://www.incompleteideas.net/Incldeas/BitterLesson.html>
- Sutton, R. S., & Barto, A. G. 2018. *Reinforcement learning: An introduction* (2nd ed.). Cambridge, MA: MIT Press.
- Sycara, K. P. 1998. Multiagent systems. *AI Magazine*, 19: 79–92.
- Tenenbaum, J. B., & Griffiths, T. L. 2001. Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, 24: 629–640.
- Tesauro, G. 1995. Temporal difference learning and TD-Gammon. *Communications of the ACM*, 38: 58–68.
- Thompson, J. 1967. *Organizations in action: Social science bases in administrative theory*. New York, NY: McGraw-Hill.
- 't Hooft, G. 2016. *The cellular automaton interpretation of quantum mechanics*. New York, NY: Springer.
- Thorndike, E. L. 1911. *Animal intelligence*. New York, NY: Macmillan.
- Turing, A. M. 1950. Computing machinery and intelligence. *Mind*, LIX: 433–460.
- Vapnik, V. N. 1995. *The nature of statistical learning theory*. New York, NY: Springer.
- Varela, F. G., Maturana, H. R., & Uribe, R. 1974. Autopoiesis: The organization of living systems, its characterization and a model. *Bio Systems*, 5: 187–196.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., Vezhnevets, A. S., Leblond, R., Pohlen, T., Dalibard, V., Budden, D., Sulsky, Y., Molloy, J., Paine, T. L., Gulcehre, C., Wang, Z., Pfaff, T., Wu, Y., Ring, R., Yogatama, D., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Kavukcuoglu, K., Hassabis, D., Apps, C., Silver, D., & Wunsh, D. 2019. Grandmaster level in *StarCraft II* using multi-agent reinforcement learning. *Nature*, 575: 350–354.
- von Krogh, G. 2018. Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Academy of Management Discoveries*, 4: 404–409.
- von Neumann, J. 1956. Probabilistic logics and the synthesis of reliable organisms from unreliable components. In C. E. Shannon & J. McCarthy (Eds.), *Automata studies*: 43–98. Princeton, NJ: Princeton University Press.
- von Neumann, J. 1958. *The computer and the brain*. New Haven, CT: Yale University Press.
- von Neumann, J., & Burks, A. W. 1966. *Theory of self-reproducing automata*. Urbana, IL: University of Illinois Press.
- von Neumann, J., & Morgenstern, O. 1944. *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

- Waldrop, M. M. 1992. *Complexity: The emerging science at the edge of order and chaos*. New York, NY: Touchstone.
- Wallach, W., & Allen, C. 2009. *Moral machines: Teaching robots right from wrong*. Oxford, U.K.: Oxford University Press.
- Walsh, J. P. 1995. Managerial and organizational cognition: Notes from a trip down memory lane. *Organization Science*, 6: 280–321.
- Wegner, D. M. 1987. Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior*, vol. 9: 185–208. New York, NY: Springer.
- Wegner, D. M. 1995. A computer network model of human transactive memory. *Social Cognition*, 13: 319–339.
- Wegner, D. M., Giuliano, T., & Hertel, P. 1985. Cognitive interdependence in close relationships. In W. J. Ickes (Ed.), *Compatible and incompatible relationships*: 253–276. New York, NY: Springer.
- Weick, K. E. 1979. *The social psychology of organizing* (2nd ed.). Reading, MA: Addison-Wesley.
- Weiss, G. (Ed.). 2013. *Multiagent systems* (2nd ed.). Cambridge, MA: MIT Press.
- Whitehead, A. N., & Russell, B. 1910. *Principia mathematica*, vol. 1. Cambridge, U.K.: University Press.
- Wiener, N. 1948. *Cybernetics*. New York, NY: J. Wiley.
- Winograd, T., & Flores, F. 1986. *Understanding computers and cognition: A new foundation for design*. Norwood, NJ: Ablex.
- Winston, P. H. 1977. *Artificial intelligence*. Reading, MA: Addison-Wesley.
- Winston, P. H. 1992. *Artificial intelligence* (3rd ed.). Reading, MA: Addison-Wesley.
- Winter, S. G. 1987. Knowledge and competence as strategic assets. In D. J. Teece (Ed.), *The competitive challenge: Strategies for industrial innovation and renewal*: 159–184. Cambridge, MA: Ballinger.
- Wolfram, S. 1983. Statistical mechanics of cellular automata. *Reviews of Modern Physics*, 55: 601–644.
- Wolfram, S. 2004. *A new kind of science*. Champaign, IL: Wolfram Media.
- Wolfram, S. 2020. *A project to find the fundamental theory of physics*. Champaign, IL: Wolfram Media.
- Wolpert, D. H. 1996. The lack of a priori distinctions between learning algorithms. *Neural Computation*, 8: 1341–1390.
- Wolpert, D. H. 2013. What the no free lunch theorems really mean: How to improve search algorithms. *Ubiquity*, 2013: 1–15.
- Wooldridge, M. J. 2009. *An introduction to multiagent systems* (2nd ed.). Chichester, U.K.: John Wiley & Sons.
- Wright, S. 1921. Correlation and causation. *Journal of Agricultural Research*, 20: 557–580.
- Wright, S. 1932. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In D. F. Jones (Ed.), *Proceedings of the sixth annual congress of genetics*: 356–366. Austin, TX: Genetics Society of America.
- Zajac, E. J., & Westphal, J. D. 2004. The social construction of market value: Institutionalization and learning perspectives on stock market reactions. *American Sociological Review*, 69: 433–457.
- Zhou, Z. H. 2012. *Ensemble methods: Foundations and algorithms*. Boca Raton, FL: Taylor & Francis.
- Zuzul, T. W. 2019. “Matter battles”: Cognitive representations, boundary objects, and the failure of collaboration in two smart cities. *Academy of Management Journal*, 62: 739–764.



Felipe A. Csaszar (fcsaszar@umich.edu) is an associate professor of strategy at the Ross School of Business, University of Michigan. He received his PhD from the Wharton School at the University of Pennsylvania. He is interested in combining formal models and empirical analyses to study issues related to the foundations of strategy and organizations.

Tom Steinberger (tomsteinberger@kaist.ac.kr) is an assistant professor at the School of Business & Technology Management, KAIST. His PhD is from the University of California, Irvine. His research is about representation and reasoning in organizational search, problem-solving, and design. He is interested in formalizing the relationship between organizational architecture and intelligence.

