

## COMBINING HUMAN AND ARTIFICIAL INTELLIGENCE: HYBRID PROBLEM-SOLVING IN ORGANIZATIONS

SEBASTIAN RAISCH  
KATERYNA FOMINA  
University of Geneva

**Organizations increasingly use artificial intelligence (AI) to solve previously unexplored problems. While routine tasks can be automated, the intricate nature of exploratory tasks, such as solving new problems, demands a hybrid approach that integrates human intelligence with AI. We argue that the outcomes of this human–AI collaboration are contingent on the processes employed to combine human intelligence and AI. Our model unpacks three hybrid problem-solving processes and their outcomes: Compared to human problem-solving, autonomous search generates more distant solutions, sequential search enables more local solutions, and interactive search promotes more recombinative ones. Collectively, these hybrid problem-solving processes broaden the range of organizational search outcomes. We enrich the behavioral theory of the firm with a technology-conscious perspective of organizational problem-solving that complements its traditional human-centric perspective. Additionally, we contribute to the literature on AI in management by extending its scope from using predictive AI for routine tasks to generative AI applications for more exploratory tasks.**

All life is problem solving.

—Karl Popper, *Philosopher of Science*

We are now solving problems with machine learning and artificial intelligence that were in the realm of science fiction for the last several decades.

—Jeff Bezos, *Chairman of Amazon*

Simon et al. (1987: 11) described the tasks of managers, scientists, and engineers in organizations as “largely [the] work of making decisions and solving problems.” Following in their footsteps, behavioral theory scholars developed models of decision-making and problem-solving in organizations that “took account of the nature of the human agents who constituted them” (Puranam, Stieglitz, Osman, & Pillutla, 2015: 337). These models describe how

humans’ cognitive limitations, and the incomplete information they process, bias organizational decisions (Gavetti, Greve, Levinthal, & Ocasio, 2012; March & Simon, 1958) and constrain organizational search for problem solutions (Cyert & March, 1963; Posen, Keil, Kim, & Meissner, 2018).

Organizations’ increasing use of artificial intelligence (AI) challenges these human-centric assumptions. AI enables artificial agents to perform cognitive functions, such as decision-making and problem-solving, previously only associated with humans (Krakowski, Luger, & Raisch, 2023). Research on AI in management has suggested that these artificial agents do not have humans’ cognitive limitations (Murray, Rhymer, & Sirmon, 2021), and that their predictions are often superior to those of humans (Agrawal, Gans, & Goldfarb, 2018: 110). Consequently, scholars have explored how, and under which conditions, AI prediction reduces biases in recurrent decisions (Shrestha, Ben-Menahem, & Von Krogh, 2019), such as having to select between candidates for positions (Newman, Fast, & Harmon, 2020), deals for venture capital investments (Blohm, Antretter, Sirén, Grichnik, & Wincent, 2020), or customer offerings for sales calls (Bader & Kaiser, 2019).

While such routine decision-making focuses on previously explored situations with known procedures and solution alternatives, problem-solving

---

The authors wish to thank associate editor Richard Makadok and three anonymous reviewers for their guidance on developing the paper. Furthermore, the authors are grateful to Tina Ambos, Oliver Baumann, Felipe Csaszar, Thomas Keil, Sebastian Krakowski, Dirk Martignoni, Jonathan Schad, Katherine Tatarinov, and the participants of research seminars at the University of Geneva and WU Vienna for their valuable comments on earlier versions of the paper. This research received support from the Swiss National Science Foundation (SNSF grants 185164 and 206856). Accepted by Richard Makadok.

occurs in situations where the problem is new and the solutions are unknown (March & Simon, 1958: 160; Winter, 2003). The principal challenge, therefore, lies not in predicting and selecting the best alternative, but in understanding a previously unexplored problem and embarking on a search for new solutions (Posen, Keil, Kim, & Meissner, 2018). Consequently, problem-solving is not limited to the use of predictive AI, which learns patterns from existing data to anticipate future outcomes, but could also benefit from generative AI, which creates new data based on learned patterns (Savage, 2023).

Furthermore, while routine tasks can be automated, the intricate nature of more exploratory tasks, such as problem-solving, demands a hybrid approach that integrates human intelligence with AI (Von Krogh, 2018). Some scholars have suggested that using AI for problem-solving might aid humans by enabling more distant organizational searches (Amabile, 2020; Raisch & Krakowski, 2021). However, others have cautioned that AI could interfere with human behavior by, for example, imposing formal rationality (Lindebaum, Vesa, & den Hond, 2020), thereby exacerbating organizations' learning myopia (Balasubramanian, Ye, & Xu, 2022). It is therefore unclear how AI technology's promise to enable more distant searches translates into actual search processes and outcomes in organizational contexts where humans and AI agents jointly solve problems.

We address this lacuna by investigating the processes and outcomes associated with combining human intelligence and AI in organizational problem-solving. Our *hybrid problem-solving* model conceptualizes three processes: *Autonomous search* combines predictive and generative AI to create solutions independently, while humans select from the solutions. For example, GM used AI to design ultra-light car parts that helped the company meet new carbon emission standards (Krok, 2018). An AI agent generated thousands of new designs and improved them through machine learning. GM's engineers selected a new AI-generated design that is 40% lighter and 20% stronger than the original version. *Sequential search* starts with predictive AI exploring a problem, but thereafter humans search for solutions. For example, scientists used an AI agent to identify the biological pathways that a Covid-19 treatment could target (Kuchler, 2022). These insights then allowed the scientists to repurpose an existing drug, which is now widely used as a Covid-19 treatment. *Interactive search* uses predictive and generative AI, but allows humans to search jointly with AI. For example, an ailing U.K. retailer's

creative team worked with AI to learn which advertising contents trigger campaign sales (Dempsey, 2021). These insights led to jointly developed new advertising contents that not only differed markedly from the company's traditional contents but also turned its sales around.

Drawing on these hybrid problem-solving types, we build theory on how each type involving a human and an AI agent changes the problem-solving process and outcome compared to an equivalent collective search with two humans. Our theory suggests that autonomous search is associated with more distant outcomes, because an AI agent searches more widely for solutions than a human can, and the human tends to follow AI's quantitative ranking of its solutions. Conversely, we propose that sequential search leads to more local outcomes, because an AI agent's analysis of a problem provides better search directions, which, in turn, increase the chances of a human's local solution search being successful. Finally, we surmise that interactive search leads to more recombinative outcomes, since a human and an AI agent's mutual learning promotes the integration of new and existing knowledge. We conclude by identifying time and expertise as moderators of the hybrid problem-solving processes' effects on outcomes.

We enrich the behavioral theory of the firm (Gavetti et al., 2012; March & Simon, 1958) with a technology-conscious perspective of organizational problem-solving that complements its traditional human-centric perspective. Our work conceptualizes several new search mechanisms that emerge from AI or human–AI interaction. It also reveals that hybrid problem-solving widens the range of organizational search outcomes compared to those that current behavioral search models predict (Cyert & March, 1963; Katila & Ahuja, 2002). In addition, our perspective explains these outcomes as being the result of the complementarities between human and AI agents' asymmetric capabilities in the problem-solving process, rather than by linking agents' capabilities directly to the search outcomes (Puranam et al., 2015; Simon, 1957). Finally, we contribute to the literature on AI in management (Murray et al., 2021; Raisch & Krakowski, 2021) by extending its scope from using predictive AI for routine tasks to generative AI applications for more exploratory tasks, such as search and problem-solving.

## ORGANIZATIONAL SEARCH THEORY

There is a long research tradition of exploring how organizations solve problems by drawing on the

behavioral theory of the firm (Katila & Ahuja, 2002; March & Simon, 1958). An organization's recognition of a problem leads to a search process that ceases once a satisfactory solution is found (Cyert & March, 1963). This search process is central to a broad variety of organizational behaviors, including the creation of novel strategies, the pursuit of entrepreneurial activities, and the development of new products (Greve, 2003). Owing to incomplete information, organizations must search for solutions, although their human agents' cognitive limitations constrain this search to a small set of alternatives (Simon, 1957). These agents pursue local search "in the neighborhood of the current alternative," meaning that a "new solution will be found 'near' an old one" (Cyert & March, 1963: 170). Only when a local search fails to find a solution does the organization gradually move toward a more distant search (Gavetti et al., 2012).

A key downside of local search is that it is less likely to generate the variability required to solve novel problems (Fleming & Sorenson, 2004). Moreover, local search may not lead to the best solutions in complex environments, which could see organizations stalling due to their inferior solutions (Levinthal, 1997). Local search's limitations have spawned a rich body of literature on the mechanisms that organizations use to promote more distant search (e.g., Gavetti & Levinthal, 2000; Knudsen & Levinthal, 2007), which include technological tools, such as knowledge repositories (Furlan, Galeazzo, & Paggiaro, 2019), crowdsourcing platforms (Afuah & Tucci, 2012), and online communities (Jeppesen & Lakhani, 2010). However, these studies have also shown that the use of technology-as-a-tool requires substantial cognitive capacity, which frequently overwhelms boundedly rational humans using such tools, and who therefore continue constraining their search to a limited set of alternatives (Afuah & Tucci, 2012; Piezunka & Dahlander, 2015).

Contrary to prior technologies, AI agents can search independently of humans (Amabile, 2020; Von Krogh, 2018). Further, unlike humans, they can process a quasi-unlimited set of alternatives (Raisch & Krakowski, 2021), and often produce better predictions of their performance (Agrawal et al., 2018: 110). Considering these technological changes, we develop a conceptual framework specifying an AI-based search's key applications to organizational problem-solving processes. We start by clarifying our theory's boundaries as a set of baseline assumptions (Dubin, 1978) referring to the characteristics of an organizational search embedded in our theorizing, as well as

to the delineation of the contexts to which our theory applies.

### Assumptions about Organizational Search

**Context.** We reaffirm that organizational search occurs under conditions of uncertainty—meaning that when organizational agents attempt to solve new problems, their knowledge about the problem (March & Simon, 1958: 161), the range of possible solutions (Kaplan, 2011), and their outcomes (Posen et al., 2018) is incomplete. This assumption corresponds to prior studies conceptualizing uncertainty as the informational context within which organizations search (e.g., Fleming, 2001; Nelson & Winter, 1982; Simon, 1957).

**Agency.** Conditions of uncertainty imply that human agents influence the search by using their cognitive capabilities. For example, humans learn from observation (i.e., vicarious learning) and experimentation (i.e., experiential learning), which enables them to apply their expertise and creativity to new problems (Gavetti & Levinthal, 2000). However, like prior research (e.g., Csaszar & Levinthal, 2016; Fleming, 2001; March & Simon, 1958), we recognize that because humans lack complete information, and have limited cognitive capacity to process this information, their search could also be constrained. Specifically, prior research has suggested that humans form mental representations that only consider certain problem dimensions, which makes their understanding of a problem partial, constrains the set of possible solutions they consider, and biases their evaluation of these options (Gavetti & Levinthal, 2000; Knudsen & Levinthal, 2007).

**Process.** We follow prior work describing the search process as comprising two stages (e.g., Csaszar & Levinthal, 2016; March & Simon, 1958: 161; Posen et al., 2018): problem definition and solution search. During problem definition, organizational agents identify the elements that constitute the problem and clarify the relationships between these elements. By doing so, the agents form a mental representation of the space in which solution search could be undertaken. During a solution search, humans explore this space to generate solutions. These activities could be connected through feedback loops, and organizations could cycle back and forth between the stages.

**Outcomes.** Finally, we assume that humans' cognitive limitations contribute to organizations' tendency to generate local search outcomes rather than more distant ones. This assumption recognizes that local search requires fewer cognitive resources,

leverages current expertise, and often leads to superior short-term performance (Knudsen & Levinthal, 2007; Laursen, 2012). Although this assumption does not always apply (Billinger, Stieglitz, & Schumacher, 2014), humans' search processes are *ceteris paribus* more likely to result in local search outcomes (e.g., Greve, 2003; Katila & Ahuja, 2002; Stuart & Podolny, 1996).

We next integrate insight from the literature on AI in management to expand our assumptions from human to artificial agents.

## AI IN MANAGEMENT

Foundational AI research has addressed managerial applications such as decision-making (Newell, Shaw, & Simon, 1959). Simon (1965: 47) predicted that "we will soon have the technological means (...) to automate all managerial decisions." While expectations were high, subsequent technological progress was slow, which led to an "AI winter" (Csaszar & Steinberger, 2021: 23). In recent years, advances in computational power, the exponential increase in data, and new machine-learning models have enabled managerial practice to adopt AI (Raisch & Krakowski, 2021). Current applications mostly use deep learning with artificial neural networks (LeCun, Bengio, & Hinton, 2015), a specific type of AI that differs from prior technologies in its unique ability to learn and act autonomously. Artificial neural networks are computing systems comprising neurons and layers that simulate the human brain structure. Deep learning is a class of machine-learning algorithms that uses artificial neural networks' multiple layers to progressively extract higher-level features from the input data (Shrestha, Krishna, & Von Krogh, 2021).

These AI technologies are now used for problem-solving in various application domains, such as drug discovery,<sup>1</sup> industrial design, and content creation.

<sup>1</sup> DeepMind's AlphaFold (Jumper et al., 2021) provides an impressive example of AI-based problem-solving in drug discovery. AlphaFold solved the "protein folding problem," which had been a grand challenge in biology for half a century. The AI-based platform predicts the biological functionality that proteins unfold when entering organisms such as the human body. Before AlphaFold, centuries of experiments and laboratory studies had led to predictions of roughly 180,000 proteins. DeepMind delivered predictions of 100 million more proteins; that is, nearly every protein whose genetic sequence is known to science. Since DeepMind open-sourced AlphaFold in 2021, research institutions have started leveraging this AI-based

Pharmaceutical companies including Pfizer, Roche, and Sanofi use them to discover drugs when there are either no drugs for a disease or the existing drugs have limited efficacy (Deng, Yang, Ojima, Samaras, & Wang, 2022; Fleming, 2018). Manufacturing companies, such as Airbus, GM, and Volkswagen, use AI-based search to design industrial products or components in situations with complex design problems for which there are either no solutions or those that do exist are insufficient (Oh, Jung, Kim, Lee, & Kang, 2019; Vinoski, 2019). Consumer product companies, such as L'Oréal, Sephora, and Unilever, use AI-based search to create audio, text, and visual content for advertising, marketing, and social media if there is either no appropriate content or the current content fails to attract sufficient consumer interest (Anantrasirichai & Bull, 2021; Simone, 2021).

While management scholars have taken note of such AI applications to solve problems (Amabile, 2020; Von Krogh, 2018), they have not explored them further. Instead, recent research on AI in management has focused on routine decision-making (e.g., Balasubramanian et al., 2022; Shrestha et al., 2019). This work provides general insight into AI applications in management, which we integrate to expand our baseline assumptions from human agents to artificial ones.

## Assumptions About AI Agents

**Context.** Building on prior research, we assume that humans must remain involved when organizations use AI to solve new problems under uncertainty (Von Krogh, 2018).<sup>2</sup> When new problems are addressed, humans need to set objectives and provide input data to allow AI agents to operate (Raisch & Krakowski, 2021). Furthermore, solving problems under uncertainty is computationally difficult; consequently, instead of optimizing these problems, AI agents just approximate them; that is, the solutions they generate might be close enough to be practically useful, but they nevertheless always relax certain real-life constraints (Fortnow, 2013). Consequently, humans need to use their intuition and judgment to

platform to discover drugs and design *de novo* protein candidates with applications in biotechnology, medicine, agriculture, food science, and bioengineering (Toews, 2021).

<sup>2</sup> This first assumption holds if there is no artificial general intelligence. Computer scientists agree widely that there will be no artificial general intelligence in the foreseeable future (Walsh, 2017).

reconcile AI agents' output with reality when selecting solutions (Brynjolfsson & McAfee, 2014: 92).

**Agency.** We further assume that AI agents' greater information-processing capacity allows them to explore a search space more widely than humans can (Amabile, 2020; Raisch & Krakowski, 2021). However, AI agents only do so when they have access to extensive data with many prior solution examples to define and explore the search space. While the digital age's increasing data generation has extended AI's range of application domains greatly (Von Krogh, 2018), insufficient data availability could still lead to suboptimal system performance or prevent the use of AI entirely (Choudhury, Starr, & Agarwal, 2020). Contrary to humans, who can use their intuition and creativity, AI cannot search outside the space emerging from the data.

**Process.** Consistent with prior research, our model covers autonomous (Balasubramanian et al., 2022), sequential (Lebovitz, Lifshitz-Assaf, & Levina, 2022), and interactive (Metcalf, Askay, & Rosenberg, 2019) processes of AI use in organizations (Shrestha et al., 2019). This variation is due to AI agents' superiority in certain dimensions, such as being able to process data more comprehensively (Murray et al., 2021). Nevertheless, humans outperform in other dimensions, such as being able to use their expertise to generate creativity (Brynjolfsson & McAfee, 2014: 202).

**Outcomes.** Finally, research has suggested that outcomes vary, depending on how human and AI capabilities are used in managerial processes (Murray et al., 2021; Raisch & Krakowski, 2021). When humans are involved in these processes, the outcomes are the direct result not of AI's capabilities but of the ways in which human and AI capabilities are used. Accordingly, we focus our theory not on AI's technological abilities but on its actual use in organizations.<sup>3</sup>

While prior research on AI in management informs our study, we next move from baseline assumptions to building a specific theory of AI use for problem-solving in organizations.

## HYBRID PROBLEM-SOLVING

Hybrid problem-solving is a search process that organizations use to solve problems by combining

<sup>3</sup> Consistent with prior research's interest in AI's use rather than its development, we further assume that AI technology is available at the start of the problem-solving process.

human and artificial intelligence. We first discuss this hybrid problem-solving process in general and then distinguish between three types of hybrid problem-solving (see Figure 1).

### The Hybrid Problem-Solving Process

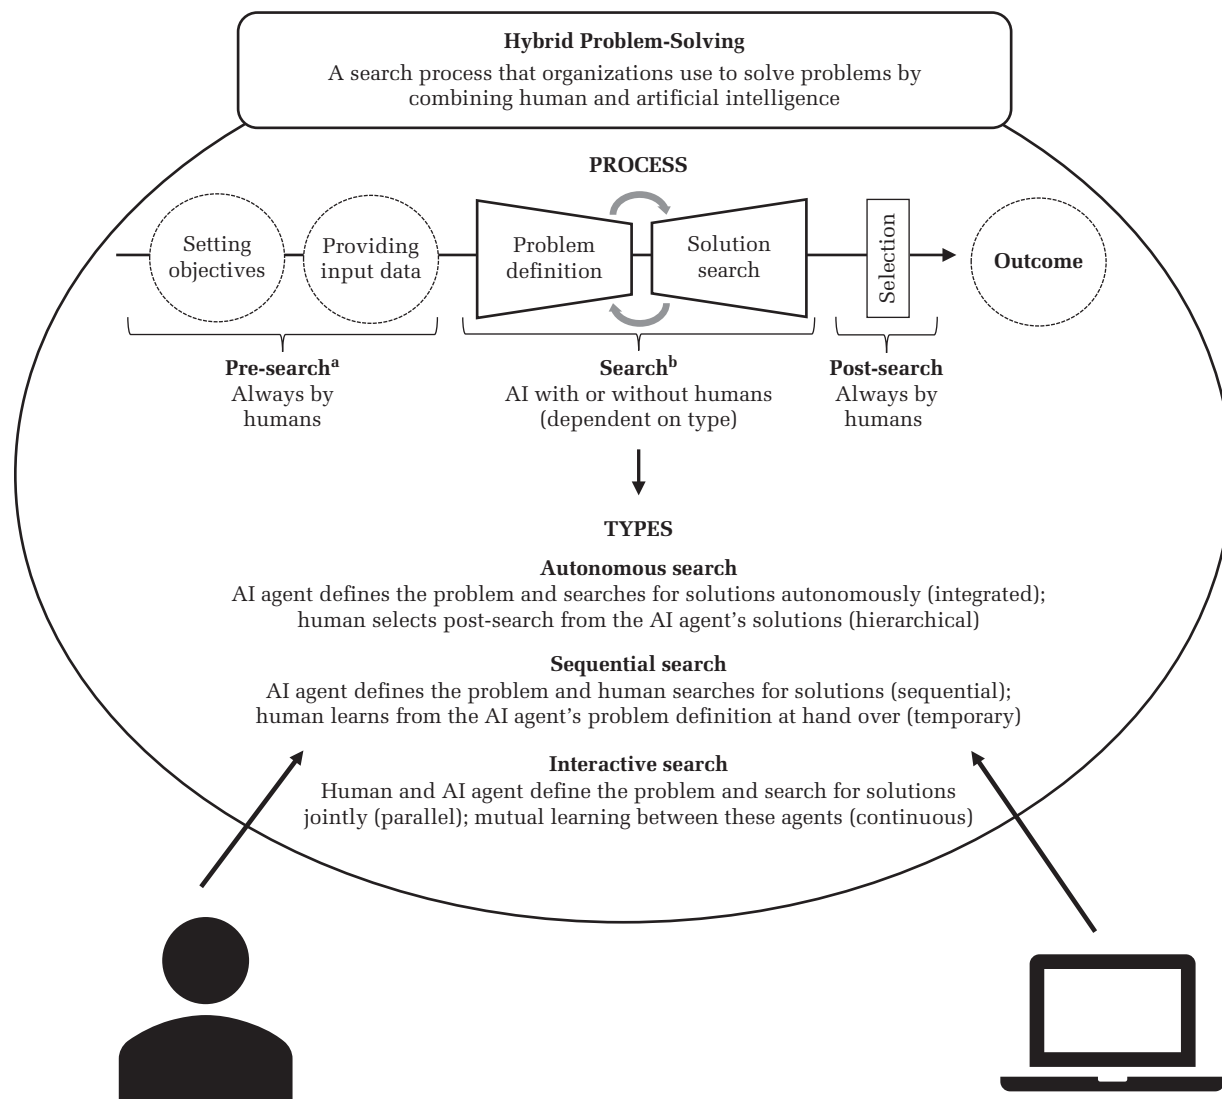
**Pre-search stage.** Humans initially engage in *setting objectives*, or specifying an AI analysis's target function (Russell & Norvig, 2020), usually with multiple requirements. For example, a new drug should demonstrate not only high therapeutic efficacy against the target disease but also, for example, high selectivity (to minimize side effects) and low toxicity. Since it is generally impossible to define a full spectrum of target function requirements, AI analysis tends to relax some of the real-life constraints (Fortnow, 2013).

Humans are also responsible for *providing input data*, usually by combining multiple open or proprietary data sources. In drug discovery, for example, pharmaceutical companies use public chemical compound libraries with descriptions of 100,000s of molecules, as well as natural language processing-generated databases of scientific publications on the target disease, drugs, and patents (Deng et al., 2022). Input data are critical for AI agents' performance. If data samples do not reflect the underlying distribution, the search results might be biased (Choudhury et al., 2020). While sufficiently objective or unbiased data are likely to be available in the application domains we discuss in this paper, this is unlikely to be universally true.

**Search stage.** There are two AI applications that organizations could use to search: predictive AI and generative AI.<sup>4</sup> In the *problem definition*, predictive AI learns a previously unexplored problem's representation directly from the input data. For example, AI predicts the molecular features, or the combinations of such features, associated with high therapeutic efficacy against the target disease, high selectivity, and low toxicity. Deep learning forms representations at multiple levels of abstraction (Bengio, 2012) by moving gradually from identifying individual features associated with a single objective (feature learning) to representing multiple interrelated features

<sup>4</sup> Some scholars have distinguished between discriminative and generative AI models (e.g., Jebara, 2004; Ng & Jordan, 2001). While discriminative models are also predictive models, the former are limited to classifying data into predefined categories while the latter also include regressions or the prediction of continuous values.

**FIGURE 1**  
**Hybrid Problem-Solving**



<sup>a</sup> Held constant in our model.

<sup>b</sup> The two search stages may be recursive.

associated with all the target objectives (representation learning) (LeCun et al., 2015). The resulting representation serves to predict the search space containing all possible solutions exhibiting the target features, called the latent space (Fernandes, Correia, & Machado, 2019).

In *solution search*, generative AI searches for new solutions exhibiting the target features in the latent space (Goodfellow et al., 2014). For example, AI discovers molecules, or combinations of molecules, with

high therapeutic efficacy against the target disease, high selectivity, and low toxicity. The latent space to be explored can be huge. For example, the chemical space for drug discovery contains  $10^{60}$  theoretically possible molecules, most of which have never been explored (Deng et al., 2022). By sampling from the latent space, generative AI creates new synthetic data and solutions beyond the input data (Tanaka & Aranha, 2019). Novelty can arise from the use of previously unexplored molecules or new molecule combinations.

While the problem definition precedes the solution search, humans (Posen et al., 2018) and AI agents (Goodfellow, Bengio, & Courville, 2016: 301) usually iterate between these stages.<sup>5</sup>

**Post-search stage.** After the search has been completed, humans are responsible for the final solution's *selection*. They use their contextual understanding to assess solutions and select from these. This selection could therefore be subject to human biases. Research has shown that humans prefer more proximate solutions (Greve, 2003; Knudsen & Levinthal, 2007) and have difficulties with assessing solutions accurately across objectives (Ethiraj & Levinthal, 2009). Predictive AI suffers less from these limitations and could therefore inform humans' selection by providing a quantitative assessment of AI-generated solutions across objectives (LeCun et al., 2015).

**Outcome.** Finally, we follow Katila and Ahuja (2002) when assessing local versus distant search across two dimensions: the *search depth*, which measures the extent to which existing solution knowledge is reused (i.e., knowledge available at the start of the problem-solving process), and the *search scope*, which indicates the extent to which new knowledge is used (i.e., knowledge generated during the problem-solving process). The reason for this distinction is that the search depth and scope are contradictory, but also mutually enabling, dimensions: Existing knowledge is not only required to absorb and integrate new knowledge (Zahra & George, 2002) but also enables recombinations that are a major source of novelty (Fleming, 2001). Distinguishing between the search depth and the search scope is also important for hybrid problem-solving, since scholars have suggested that, given its superior information-processing capacity, AI has a greater search scope (Von Krogh, 2018), while humans, who have richer expertise, have a greater search depth (Brynjolfsson & McAfee, 2014: 202). While the search depth and scope vary during the search process, we assess these dimensions in terms of the solution that organizations select for implementation.<sup>6</sup>

<sup>5</sup> While we consider iterations between the two search stages, we assume linear progression in the rest of the process.

<sup>6</sup> We focus on the solution that organizations select for implementation due to its impact on more distant outcomes, such as the solution's performance. We do not assess the solution's performance, since factors outside our theory-building efforts' scope, such as the organizational implementation and the market conditions, also impact it.

## Types of Hybrid Problem-Solving

Hybrid problem-solving requires at least two agents—one human and one artificial—to address a problem. While prior search studies have generally focused on an individual agent's problem-solving, some scholars have recognized that complex problems create information-processing demands that often exceed any individual agent's cognitive capacity (e.g., Baumann, 2015; Levinthal & Posen, 2007). In such situations, two or more agents should address the problem collectively.

A key collective problem-solving challenge is that organizational tasks are usually not perfectly decomposable, which creates interdependencies between agents (Heath & Staudenmayer, 2000; Simon, 1962). Prior work has described three types of search task division and their respective coordination mechanisms (Billinger, Benincasa, Baumann, Kretschmer, & Schumacher, 2023; Cyert & March, 1963: 200). The first type requires agents to conduct the entire search task separately (integrated search), while a higher-level agent with contextual understanding selects from their solutions (hierarchical coordination) (Knudsen & Levinthal, 2007; Rivkin & Siggelkow, 2003). The second type makes agents conduct different search subtasks sequentially (sequential search), thereby allowing one-sided learning when agents integrate the previous agents' insights (temporary coordination) (Baumann, 2015). In the third type, agents work jointly on the entire search task (parallel search), engaging in mutual learning (continuous coordination) (Knudsen & Srikanth, 2014; Puranam & Swamy, 2016).

Following extant work (e.g., Billinger et al., 2023), we focus on a simple dyad of agents, in our case one human and one artificial, to develop a typology of hybrid problem-solving.<sup>7</sup> *Autonomous search* is the first hybrid problem-solving type, which requires an AI agent to conduct the entire search task (i.e., problem definition and solution search) by using predictive and generative AI in a closed loop (integrated search). Generative AI creates new solutions that are added to the input data, allowing predictive AI to update the problem definition, which, in turn,

<sup>7</sup> Our propositions could be extended to systems with more than two agents. For example, in practice, there are some AI applications that involve many humans (i.e., artificial swarm intelligence; see Metcalf et al., 2019) or AI agents (i.e., distributed AI; see Bond & Gasser, 2014).

enables the next round of solution search.<sup>8</sup> A human with contextual understanding finally selects from the AI-generated solutions (hierarchical coordination), although this person has little insight into AI's underlying black-box models (Linardatos, Papastefanopoulos, & Kotsiantis, 2020). Autonomous search is used to ensure that humans, and their biases, do not affect the search process. For example, Insilico Medicine used autonomous search to discover a drug for pulmonary fibrosis (Hale, 2021a), NASA to lighten its next-generation space suits (Oberhaus, 2020), and L'Oréal to create new social media content that doubled its returns (Prosser, 2021).

A *sequential search* uses predictive AI for the problem definition, but a human subsequently conducts a solution search without the use of generative AI (sequential search).<sup>9</sup> Coordination at handover allows the human to learn from the AI agent's problem definition (temporary coordination). Such learning is possible because sequential search uses "explainable AI" (Senoner, Netland, & Feuerriegel, 2022), which pertains to the supplementary application of interpretability methods to enhance the transparency of AI models. Sequential search is used to benefit from AI's superior prediction (Agrawal et al., 2018: 110) before introducing unique human capabilities, such as their creativity and contextual understanding (Brynjolfsson & McAfee, 2014: 92), to overcome AI's limitations. For example, BenevolentAI used sequential search to discover a Covid-19 drug (Metz, 2020), Tommy Hilfiger to create new fashion designs (Arthur, 2018), and Utah's ski resorts

to generate social media content that increased their customers' engagement (Cortex, 2022).

The final hybrid problem-solving type is *interactive search*. The human and the AI agent work jointly on the problem definition and the solution search (parallel search). Coordination throughout the search process enables the human and the AI agent's mutual learning (continuous coordination). While the inclusion of explainable AI methods enables this learning, the use of generative AI always limits the resulting models' transparency (Linardatos et al., 2020).<sup>10</sup> Interactive approaches are used to combine the human and the AI agent's complementary learning skills (Puranam, 2021). Interactive search allowed BenevolentAI to discover the first cure for a rare childhood brain cancer (Gregory, 2021), Philippe Starck to design a mass-market chair using a minimal amount of material (Schwab, 2019), and the electro-pop band Yacht to compose its first Grammy-nominated album (Chow, 2020).

In the following sections, we argue that, despite their basic similarities, there are two important differences between these hybrid types and collective problem-solving. First, collective problem-solving divides the search task between human agents with relatively homogenous cognitive abilities (Billinger et al., 2023). Consequently, prior collective problem-solving studies have not explored "agents with asymmetric abilities" (Knudsen & Srikanth, 2014: 433). However, human and AI agents have fundamentally different cognitive capabilities (Amabile, 2020; Von Krogh, 2018). Hybrid problem-solving processes and outcomes are therefore likely to vary depending on how the search task is divided between the human and the AI agent.

Second, collective problem-solving relies on coordination between humans whose similar cognitive abilities and shared understanding lead to "joint myopia" (Knudsen & Srikanth, 2014: 409). Human-AI coordination differs, because the human and the AI agent's asymmetric skills could enable complementarities, which foster learning (Choudhury et al., 2020; Puranam, 2021). However, AI opacity could constrain such learning, or even prevent it entirely (Lebovitz et al., 2022). Hybrid problem-solving processes and outcomes are therefore likely to vary with

---

<sup>8</sup> Autonomous search generally uses an AI architecture called generative adversarial networks (GANs) (Goodfellow et al., 2014). In GANs, two artificial neural networks—one for prediction and the other for generation—work together in a closed loop by means of reinforcement learning, thereby enabling continuous improvement. The generator creates possible solutions with each iteration and the predictor assesses their value against the target function. In practice, there are usually hundreds of iterations, which lead to increasingly better-performing solutions.

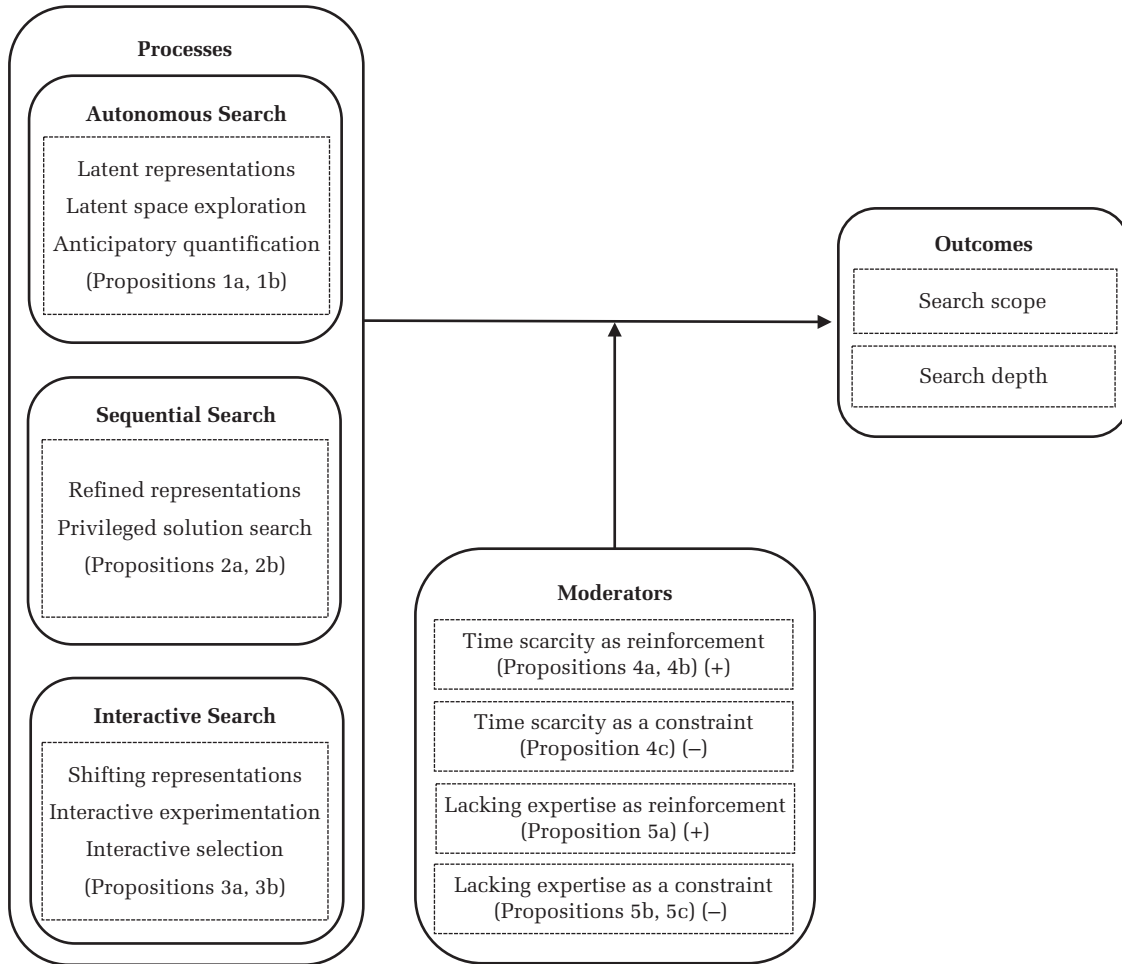
<sup>9</sup> Sequential search relies on a single artificial neural network for prediction. The type of artificial neural network used depends on the input data: If the data being processed have multiple dimensions (e.g., images), convolutional neural networks (CNNs) are an often-used type of artificial neural network, whereas recurrent neural networks are often the first choice for the sequential recognition of inputs (e.g., music); however, there are also combinations, such as 3D-CNNs, that extract both spatial and temporal features (e.g., video) (see Deng et al., 2022).

---

<sup>10</sup> While interactive search generally uses explainable GANs (Linardatos et al., 2020), which combine predictive and generative artificial neural networks, as well as supplementary interpretability methods to explain their results, generative AI models are never perfectly transparent, which limits humans' understanding of AI outputs.



**FIGURE 2**  
**Hybrid Problem-Solving: Processes, Moderators, and Outcomes**



different types of human–AI coordination, as well as the extent and nature of the learning they afford.

We next discuss the differences between autonomous, sequential, and interactive search. To make these types comparable, we hold the pre-search stage constant: The three hybrid types are assumed to address the same type of problem, with similar objectives and input data. For example, pharmaceutical companies recently used all three hybrid types to discover new drugs.<sup>11</sup> We discuss the differences between each hybrid type and a comparable collective human search for the problem-solving process

(i.e., the problem definition, solution search, and selection) and outcome (i.e., the search scope and depth). Figure 2 shows our propositions, and Table 1 provides an overview of the different problem-solving types’ processes and outcomes.

## AUTONOMOUS SEARCH

### Autonomous Search Process

**Problem definition.** In collective human search, which serves as our baseline, a human uses existing solution knowledge to form mental representations (Posen et al., 2018). Given this agent’s cognitive limitations, the resulting mental representations are partial and path-dependent (Knudsen & Levinthal, 2007). In contrast, autonomous search relies on predictive AI to form *latent representations*

<sup>11</sup> While we hold the pre-search phase constant in our theory-building efforts, we use real-life examples that inevitably exhibit some degree of variation. These examples are meant to *illustrate*, rather than replace, our theory.

TABLE 1  
Types of Hybrid Problem-Solving: Processes and Outcomes

Types	Processes			Outcomes	
	Problem Definition	Solution Search	Selection	Scope	Depth
<b>Collective human search</b> (baseline)	<ul style="list-style-type: none"> <li>Human search to form mental representations</li> <li>Mental representations are partial and path-dependent</li> </ul>	<ul style="list-style-type: none"> <li>Human search for solutions in the proximity of the existing solution knowledge</li> <li>Local search only broadens in case of failure</li> </ul>	<ul style="list-style-type: none"> <li>Human selection from solutions</li> <li>Human selection is biased toward local solutions</li> </ul>	Low	Moderate
<b>Autonomous search</b>	<ul style="list-style-type: none"> <li>AI-based search to form <i>latent representations</i> (1)</li> <li>Latent representations are more complete and less path-dependent</li> </ul>	<ul style="list-style-type: none"> <li>AI-based search to conduct <i>latent space exploration</i> (4)</li> <li>Latent space exploration is wider and more exhaustive</li> </ul>	<ul style="list-style-type: none"> <li>Human-AI coordination provides <i>anticipatory quantification</i> (7)</li> <li>AI's anticipatory quantification reduces humans' selection bias</li> </ul>	Increasing	Decreasing
<b>Sequential search</b>	<ul style="list-style-type: none"> <li>Human-AI coordination enables human learning, leading to <i>refined representations</i> (2)</li> <li>Refined representations are more complete</li> </ul>	<ul style="list-style-type: none"> <li>Human use of refined representations to conduct a <i>privileged solution search</i> (5)</li> <li>Privileged solution search increases the likelihood of the human's initial local search being successful</li> </ul>	<ul style="list-style-type: none"> <li>Human selection from solutions</li> <li>Human selection is biased toward local solutions</li> </ul>	Decreasing	Increasing
<b>Interactive search</b>	<ul style="list-style-type: none"> <li>Human-AI coordination enables mutual learning, leading to <i>shifting representations</i> (3)</li> <li>Shifting representations are more complete and less path-dependent</li> </ul>	<ul style="list-style-type: none"> <li>Human-AI coordination enables <i>interactive experimentation</i> (6)</li> <li>Interactive experimentation gradually moves to more distant areas of the search space over time</li> </ul>	<ul style="list-style-type: none"> <li>Human-AI coordination enables <i>interactive selection</i> (8)</li> <li>AI's anticipatory quantification counterbalances humans' selection bias</li> </ul>	Increasing	Similar

Note: Compared to the baseline of collective human search.

(Goodfellow et al., 2016: 528), which are less partial and path-dependent. They are less partial because the AI agent processes extensive data to progressively move from simple to highly complex representations (LeCun et al., 2015), which enable a more accurate prediction of the latent space (Bengio, Courville, & Vincent, 2013). The representations are less path-dependent because the AI agent learns them directly from the input data without first receiving human training (Bengio et al., 2013) and regardless of the organization's prior sensemaking (Von Krogh, 2018). Furthermore, the additional use of generative AI creates new synthetic data by sampling from the latent space (Goodfellow et al., 2014). Constantly adding these synthetic but plausible data samples makes the latent representations progressively more complete (Tanaka & Aranha, 2019) and less path-dependent (Shorten & Khoshgoftaar, 2019). Insilico Medicine's AI agent, for example, processed millions of documents to form latent representations of pulmonary fibrosis, a lifelong lung disease with limited treatment options. The resulting latent representation uncovered 20 previously unknown pathways that a new drug could target (Hale, 2021a).

**Solution search.** In collective human search, mental representations lead the human to conduct a local search, which ceases once a satisficing solution has been found (Cyert & March, 1963). In contrast, the AI agent conducts *latent space exploration* (Fernandes et al., 2019), which allows a wider and more exhaustive search. This search is wider because generative AI searches for new solutions across the entire latent space. Generative AI creates random solutions, while predictive AI provides feedback on how well these solutions match the latent representations (Goodfellow et al., 2014).<sup>12</sup> In an iterative process, the AI generator uses the AI predictor's feedback to create increasingly better-performing solutions. This search is more exhaustive, because the AI agent has quasi-unlimited information-processing capacity, allowing it to explore a greater variety of solutions from the latent space at low cost and time requirements (Raisch & Krakowski, 2021). This is possible due to the AI agent doing an "offline search" (Gavetti & Levinthal, 2000: 114) by predicting newly generated solutions' performance against the target function without actually implementing them. Insilico Medicine's AI agent, for example, generated 80

previously unexplored molecules and predicted their effectiveness in targeting the most promising drug pathway it had identified in the latent representations (Hale, 2021b).

**Selection.** In collective human search, selection is biased, since the human evaluator systematically chooses local solutions over more distant ones (Knudsen & Levinthal, 2007). This selection bias is likely to be less pronounced in autonomous search, because the AI agent provides the human with a complex quantitative evaluation of the generated solutions' fit across multiple objectives. Prior studies have found that such an *anticipatory quantification* reduces human discretion (Faraj, Pachidi, & Sayegh, 2018). While humans sometimes ignore AI recommendations due to algorithm aversion (Dietvorst, Simmons, & Massey, 2015), they tend to follow them closely when complex outputs create cognitive overload (Allen & Choudhury, 2022). Accordingly, Rivkin and Siggelkow (2003: 308) showed that human evaluators simply "rubberstamp" proposed solutions if the searchers have high search skills and provide the evaluators with little information. This is true of autonomous search, because the AI agent has high search skills (Von Krogh, 2018) and the evaluator finds the reasoning behind the black-box AI models' solutions opaque (Linardatos et al., 2020). Insilico Medicine's chief scientist, for example, readily nominated the best-performing molecule that the AI agent had generated as the clinical trial candidate (Hale, 2021a).

### Autonomous Search Outcome

**Search scope.** Whereas organizations using collective human search prioritize local search (Knudsen & Srikanth, 2014), autonomous search enables them to form less path-dependent latent representations and to explore the latent space more widely. Such a comprehensive search of the search space results in solutions that are, on average, more distant from those explored previously (Fleming, 2001; Schilling & Green, 2011). Moreover, Gavetti and Levinthal (2000) suggested that an offline search's lower cost and time requirements reduce the experimenting risk, which promotes more distant search. Finally, AI's anticipatory quantification (Faraj et al., 2018), and the evaluator's limited insight into AI's integrated search process (Linardatos et al., 2020), should reduce the human tendency to systematically select local solutions rather than distant ones. We therefore suggest that autonomous search leads to outcomes that, on average, include more new knowledge than those that

<sup>12</sup> Some applications use genetic algorithms (Goldberg, 1989), which work with metaheuristics inspired by the evolutionary process of natural selection (with crossover and mutation as genetic operators), instead of randomization.

collective human search produces. Consistent with this reasoning, Insilico Medicine's fibrosis drug is based on "a novel molecule" affecting "a biological target (...) that has never been tried before" (Hale, 2021b).

*Proposition 1a. Autonomous search increases the search outcomes' average scope compared to that of collective human search.*

**Search depth.** AI forms latent representations that divert from organizations' mental representations, which means that they negate prior sensemaking *ex ante* (Csaszar & Levinthal, 2016). Consequently, the new knowledge that AI generates is, on average, more distant from the existing solution knowledge than when humans search. This distance makes it difficult for humans to subsequently use their mental representations (Schilling & Green, 2011) to assess the AI's solutions. Burrell (2016: 10) concluded that "when a computer learns and subsequently builds its own representation (...) it does so without regard for human comprehension." Consequently, humans have few opportunities to apply existing solution knowledge *ex post*. We therefore expect that, on average, the final solutions derived from an autonomous search rely less on existing solution knowledge than those that collective human search identifies. Accordingly, Insilico Medicine's scientists did not use their expertise to develop the *de novo* drug for pulmonary fibrosis, which targets the disease on the basis of a previously unknown pathway (Hale, 2021a).

*Proposition 1b. Autonomous search decreases the search outcomes' average depth compared to that of collective human search.*

## SEQUENTIAL SEARCH

### Sequential Search Process

**Problem definition.** Like autonomous search, sequential search starts with predictive AI forming latent representations. However, temporary coordination allows the human to learn from the AI agent's problem definition. This learning is possible because sequential search includes explainable AI methods (Linardatos et al., 2020), which grant insight into the latent representations. Consulting these white-box models allows humans to explore more predictors and higher-order patterns than would have been possible without AI's use (Faraj et al., 2018). This learning leads to *refined representations*, which are

more complete than humans' traditional mental representations, therefore allowing for a more accurate latent space prediction (Csaszar & Levinthal, 2016). While refined representations are more complete, they are nevertheless path-dependent for the following reasons: First, unlike autonomous search, sequential search is limited to predictive AI, which means that no newly generated data are added to the input data. Second, humans learn by *selectively* integrating new knowledge from the AI agent's latent representations into the existing solution knowledge, which refines their mental representations in a path-dependent process (Posen & Levinthal, 2012) rather than changing them entirely (Tripsas & Gavetti, 2000). At BenevolentAI, for example, the AI agent visualized extensive information from the scientific literature in a knowledge graph showing the pathways that the Covid-19 virus uses to infect humans. While the knowledge graph contained thousands of relationships, the scientist analyzing it used his expertise to quickly zero in on two familiar pathways that "leapt out at him" (Metz, 2020).

**Solution search.** Refined representations enable a human to conduct a *privileged solution search*. Compared to humans' traditional solution search, privileged solution search is more likely to be local, because refined representations provide more local starting points for the search. Given humans' preference for local search (Cyert & March, 1963), local starting points' greater availability increases the odds that a human will start the solution search in the proximity of existing knowledge. Furthermore, this initial local search is more likely to be successful, because refined representations are more complete, meaning that they allow more accurate prediction (Gary & Wood, 2011), which increases the chances of finding a satisficing solution through local search (Gavetti & Levinthal, 2000). Accordingly, Csaszar and Levinthal (2016: 2042) argued that "a local search has more opportunities to get stuck in the nooks and crannies of a more elaborate representation." BenevolentAI's scientist initially searched for known drugs that could inhibit the two familiar pathways he had identified on the basis of the knowledge graph. Since this local search yielded 47 known drugs with both inhibiting effects, he stopped searching further (Metz, 2020).

**Selection.** As in collective human search, a human selects from the final solutions. Humans tend to systematically select local solutions rather than more distant ones (Knudsen & Levinthal, 2007). At BenevolentAI, for example, the scientist focused

exclusively on drugs that were already approved for medical use, and finally selected a well-known molecule (Kuchler, 2022).

### Sequential Search Outcome

**Search scope.** Refined representations offer more local starting points for humans' solution search. Furthermore, they enable better prediction, which increases the likelihood that humans' initial local search will be successful. Humans interpret such positive feedback on their initial experimentation as a sign of the current search region's munificence, which limits their further search to the vicinity of their initial search (Billinger, Srikanth, Stieglitz, & Schumacher, 2021). Sequential search should therefore help organizations identify local solutions in a greater share of their search initiatives, including some where humans would traditionally have been unable to identify a solution through a local search. In these initiatives, it is no longer necessary to search more widely (Posen & Martignoni, 2018), which should decrease the average search scope across the initiatives. Since the search remains local more often, and is terminated more rapidly, less new knowledge will be generated. In the Covid-19 drug example, BenevolentAI's scientist did not explore any new molecules, since his refined representation allowed him to rapidly and successfully identify a known molecule that inhibits the two targeted pathways (Metz, 2020).

*Proposition 2a. Sequential search decreases the search outcomes' average scope compared to that of collective human search.*

**Search depth.** Refined representations increase the chances of a human starting the solution search at familiar points and subsequently continuing this search in the existing solution knowledge's vicinity. Such a local search provides rich opportunities for reusing existing knowledge when developing solutions (Ott, Eisenhardt, & Bingham, 2017). We therefore expect that, on average, sequential search leads to outcomes that integrate more existing solution knowledge than collective human search does. This is because human search is more prone to failure, and therefore more often triggers distant search, leading to outcomes that build less on existing knowledge (Gavetti et al., 2012). Consistent with our argumentation, the Covid-19 drug that BenevolentAI identified through its sequential search is a repurposed rheumatoid arthritis drug. Scientists had "spent years exploring its effect on other viruses," which provided extensive

prior solution knowledge that was leveraged for the Covid-19 drug (Metz, 2020).

*Proposition 2b. Sequential search increases the search outcomes' average depth compared to that of collective human search.*

## INTERACTIVE SEARCH

### Interactive Search Process

**Problem definition.** Like sequential search, interactive search allows a human to initially learn from predictive AI's problem definition. However, unlike in sequential search, this person provides the AI agent with feedback, which triggers cycles of mutual learning (Holzinger, 2016). Each interaction cycle leads to small changes in the representations (Csaszar & Levinthal, 2016). Over time, these *shifting representations* evolve to become progressively more complete and less path-dependent than traditional mental representations. This is due to mutual learning often being cut short when humans interact, because their similar cognitive capabilities and rich communication promote shared knowledge and understanding, which rapidly limit variance (Knudsen & Srikanth, 2014). Such early convergence is less likely when a human and an AI agent interact, because their asymmetric cognitive capabilities promote greater variance (Von Krogh, 2018) and AI's remaining opacity limits these agents' communication to a greater extent than the rich communication between humans does (Murray et al., 2021). Such limited observability and communication act as partial isolation mechanisms that prevent the two searchers from converging quickly (Baumann, Schmidt, & Stieglitz, 2019; Fang, Lee, & Schilling, 2010). Consequently, each interaction cycle enables new variance and learning, which are an "important form of adaptation" (Gavetti & Levinthal, 2000: 127). BenevolentAI's scientist, for example, interacted with an AI agent to develop representations of a childhood brain cancer that was lacking treatment. The mutual learning allowed the scientist to eventually realize that, to be effective, a possible cure would have to target multiple pathways simultaneously (Gregory, 2021).

**Solution search.** The human and the AI agent also engage in *interactive experimentation*, which differs from traditional solution search in two ways: First, the use of generative AI to create new solutions pushes the human beyond their local search in each interaction cycle. While humans prefer more local solutions (Cyert & March, 1963), the AI agent also generates more distant ones from the latent space.

Since the AI agent ranks the solutions in terms of their performance against the target function, it is difficult for the human to disregard superior, but more distant, solutions completely (Von Krogh, 2018). Second, human feedback introduces new variance in each cycle, thereby enabling further AI generation in the next cycle. Such interactive experimentation has the potential to broaden a search sequentially, leading to more distant solutions over time (Baumann et al., 2019). The human and the AI agent's asymmetric cognitive capabilities promote this generative effect by introducing new variance in each cycle, which should increase the potential for more distant search (Posen & Martignoni, 2018). At BenevolentAI, for example, interactive experimentation allowed for exploring entirely new treatment combinations between molecules that "would not have been obvious to people" (Gregory, 2021).

**Selection.** As in autonomous search, the AI agent provides a quantitative evaluation of the generated solutions. The difference is that the human and the AI agent engage in *interactive selection*, which combines the AI agent's anticipatory quantification with the human's heuristic selection. We surmise that the AI assessment counterbalances a human's selection bias (Faraj et al., 2018). Given a human's involvement in the entire search process, this person is likely to form independent opinions of and preferences for jointly developed solutions, but is unlikely to disregard AI's anticipatory quantification completely (Von Krogh, 2018). BenevolentAI's scientist, for example, selected a highly ranked novel drug combination for clinical trials (Gregory, 2021).

### Interactive Search Outcome

**Search scope.** Shifting representations and interactive experimentation broaden a search sequentially, despite each interaction only making limited changes (Baumann et al., 2019). Since human-AI communication is more limited than the rich communication between humans (Murray et al., 2021), there is less risk of early convergence (Knudsen & Srikanth, 2014). Furthermore, these agents' asymmetric cognitive capabilities are more likely to introduce variance in each cycle, which allows further experimentation (Kaplan, 2011). Interactive experimentation between the human and the AI agent should therefore better maintain variation over time, which leads to solutions that, on average, incorporate more new knowledge generated during the problem-solving process than those that collective human search produces. At BenevolentAI, for

example, the interactive search incorporated previously unknown and jointly derived insights into "a new drug regime," resulting in a "new treatment combination" (Gregory, 2021).

*Proposition 3a. Interactive search increases the search outcomes' average scope compared to that of collective human search.*

**Search depth.** While interactive search generates more new knowledge, the human's strong involvement throughout the entire problem-solving process ensures that this new knowledge is integrated with existing knowledge. In each interaction cycle, the human relies on existing solution knowledge to absorb new knowledge (Gavetti & Levinthal, 2000). Shifting representations require the human to frequently revisit existing solution knowledge (Posen & Martignoni, 2018) and interactive experimentation's gradual learning process compels this person to integrate new and existing knowledge (Holzinger, 2016). However, the human's cognitive limitations constrain the degree to which this knowledge is integrated into solutions (Cyert & March, 1963). Similarly, research on AI in management has shown that an increasing information (Luo, Qin, Fang, & Qu, 2021) and cognitive (You, Yang, & Li, 2022) load hampers human-AI collaboration. We therefore expect that interactive search, similar to collective human search, leads to solutions that, on average, integrate moderate levels of existing knowledge. For example, the novel cancer drug combination resulting from BenevolentAI's interactive search integrates two known molecules that were "already approved to treat other types of cancer" (Gregory, 2021).

*Proposition 3b. Interactive search and collective human search are related to similar levels of search depth.*

## TIME AND EXPERTISE AS CONTINGENCY FACTORS

### The Moderating Role of Time

Prior research has suggested that time imposes search constraints (Baumann et al., 2019; Greve, 2003). AI agents substituting humans in the search process could help organizations overcome some of these constraints, since their superior information-processing capacity allows them to complete tasks more rapidly (Gregory, Henfridsson, Kaganer, & Kyriakou, 2021). However, hybrid problem-solving keeps humans in the process, which means that further exploration is needed to assess how time

scarcity moderates hybrid problem-solving types' outcomes.

**Time scarcity as reinforcement.** Time scarcity is likely to have little adverse effect on *autonomous search*, because, contrary to human search, the AI agent's information processing requires little time (Gregory et al., 2021). Insilico Medicine, for example, produced a novel pulmonary fibrosis medicine in 18 months, compared to the three to six years that the traditional process requires (Hale, 2021a). However, time constraints complicate humans' search (Greve, 2003), which should decrease the human's ability to challenge and revise the technology's outputs further (Orlikowski & Scott, 2014). Time scarcity therefore reduces human selection's constraining effect on the search scope and limits the human's ability to apply existing solution knowledge *ex post*, which decreases the search depth further. On this basis, we propose:

*Proposition 4a. Time scarcity reinforces an autonomous search's positive effect on the search scope (Proposition 1a) and its negative effect on the search depth (Proposition 1b).*

In *sequential search*, the human invests more time in refining representations, which, *ceteris paribus*, reduces the time available for a solution search (Csaszar & Levinthal, 2016). Such time constraints reduce the solution search's scope, particularly when the representations are more complex (Baumann et al., 2019). Under conditions of time scarcity, the human is therefore likely to shorten their solution search (Uotila, Keil, & Maula, 2017), which increases the probability of the resulting solutions being in the vicinity of the existing solution knowledge. The pressure that BenevolentAI experienced to rapidly find a cure for Covid-19 was one reason for limiting the solution search to existing drugs (Metz, 2020). We therefore propose:

*Proposition 4b. Time scarcity reinforces a sequential search's negative effect on the search scope (Proposition 2a) and its positive effect on the search depth (Proposition 2b).*

**Time scarcity as a constraint.** Contrary to autonomous and sequential searches, whose outcomes are reinforced, time scarcity constrains *interactive search*. Since time scarcity imposes constraints on the number of sequential trials in the search process (Uotila et al., 2017), the interactive experimentation is reduced and, therefore, is less likely to lead to more a distant search over time (Holzinger, 2016). Investing more time in the problem definition leaves less time

for a solution search, which increases the risk of becoming snagged during a local search (Csaszar & Levinthal, 2016). Alternatively, reducing the time spent on a problem definition increases the risk of human's representations being too distant from those of the AI agent, which could lead to "mutual confusion" (Knudsen & Srikanth, 2014: 409). Whatever the case, time scarcity constrains interactive search's effects by decreasing its search scope. Formally, we propose:<sup>13</sup>

*Proposition 4c. Time scarcity constrains an interactive search's positive effect on the search scope (Proposition 3a).*

### The Moderating Role of Expertise

Prior research has also highlighted that expertise—the skills and knowledge accumulated in a domain through prior learning (Choudhury et al., 2020)—is an enabler and constrainer in the search process (Cyert & March, 1963; Puranam et al., 2015). Predictive AI could be a substitute for humans' expertise (Agrawal et al., 2018), possibly reducing a lack of expertise's constraining effects. However, human expertise could also complement AI (Choudhury et al., 2020).

**A lack of expertise as reinforcement.** In *autonomous search*, a lack of expertise is unlikely to affect the search process negatively because the AI agent acts as a substitute of the human. The AI agent learns directly from the data without requiring access to human expertise (Russell & Norvig, 2020: 651). However, prior theory has suggested that inexperienced humans suffer even more from their limited understanding of AI outputs (Kellogg, Valentine, & Christin, 2020). Accordingly, Anthony (2021) found that inexperienced humans generally accept AI outputs without questioning them. A lack of expertise could therefore reduce human selection's constraining effect on the search scope even further. Consequently, we propose:

*Proposition 5a. A lack of expertise reinforces an autonomous search's positive effect on the search scope (Proposition 1a) and its negative effect on the search depth (Proposition 1b).*

**A lack of expertise as a constraint.** Conversely, a lack of expertise is likely to constrain *sequential search*, because the more complex representations

<sup>13</sup> We do not convey moderating propositions for interactive search's relationship with search depth, since the main relationship (Proposition 3b) predicts no difference compared to a collective human search.

that it informs swiftly overwhelm humans with little expertise (Csaszar & Ostler, 2020). Further, an inexperienced human has a lower absorptive capacity to grasp the complexities arising from AI-based technologies' use (Choudhury et al., 2020). Consequently, this human's learning from AI could be more limited. We therefore expect a lack of expertise to constrain sequential search by leading to less refined representations, which provide fewer and less accurate local starting destinations. This increases the risk of humans' initial local search being unsuccessful, which could broaden their search gradually (Gavetti et al., 2012). Furthermore, a lack of expertise constrains the human's ability to leverage prior solution knowledge fully, which could in turn reduce the search depth. Formally, we propose:

*Proposition 5b. A lack of expertise constrains a sequential search's negative effect on the search scope (Proposition 2a) and its positive effect on the search depth (Proposition 2b).*

There are similar constraints in *interactive search* when an inexperienced human's limited ability to absorb and use the AI agent's outputs, and to provide meaningful feedback, undermine the human and the AI agent's capacity to learn from one another, which could lead to "mutual confusion" (Knudsen & Srikanth, 2014: 409). A lack of expertise constrains interactive search by undermining the joint representation learning and interactive experimentation, which enable a gradual increase in the search scope. These arguments lead to our final proposition:

*Proposition 5c. A lack of expertise constrains an interactive search's positive effect on the search scope (Proposition 3a).*

## DISCUSSION

In the behavioral tradition, prior research advanced a human-centric perspective of organizational search (March & Simon, 1958; Puranam et al., 2015). We complement this work with a technology-conscious perspective explaining how organizations' increasing use of AI-enabled search for problem-solving changes the search processes and outcomes.

### Theoretical Implications

Our work's first contribution is the conceptualization of eight new AI-enabled search mechanisms (see Table 1, numbered from 1 to 8) with theoretical implications that cannot be derived from the current behavioral search models with a human-centric perspective:

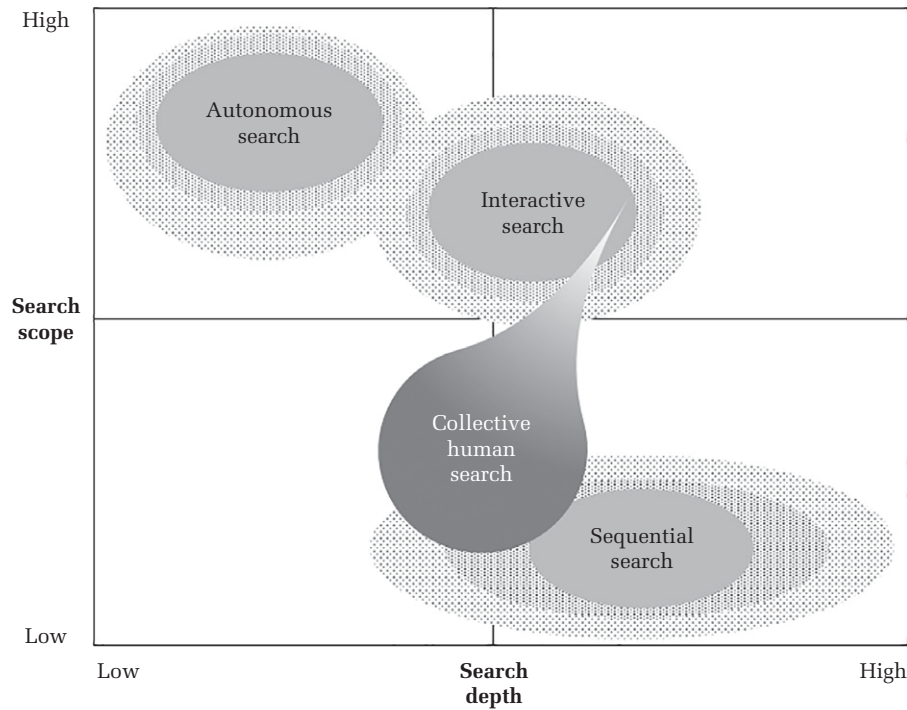
First, predictive AI allows organizations to form more complex representations than the simplistic ones that the human-centric perspective describes (Barr, Stimpert, & Huff, 1992; Csaszar & Levinthal, 2016). Three new problem definition mechanisms make this possible: Predictive AI forms highly complex and opaque latent representations (1), while combining human and AI agents' predictions creates refined (2) or shifting (3) representations that are more moderately complex, but are explainable. These AI-enabled representations can allow for a better prediction of the latent space than prior accounts of human-derived representations assumed possible (Gavetti & Levinthal, 2000). Furthermore, better prediction enables a privileged solution search (5), which is a new human solution search mechanism that benefits from AI's superior problem definition.

Second, generative AI allows organizations to conduct a broader solution search than has been described in prior research (Cyert & March, 1963; Posen et al., 2018). Latent space exploration (4) is the first solution search mechanism that is independent of human search, and therefore less affected by humans' search limitations than prior technological tools intended to broaden solution search (Afuah & Tucci, 2012; Piezunka & Dahlander, 2015). This new AI-based search mechanism reduces organizations' data-availability limitations (Simon, 1957) by generating new data, their data-processing limitations (Cyert & March, 1963) by using machines' extensive capacity, and their path dependencies (Gavetti & Levinthal, 2000) by learning directly from the data. The other mechanism, interactive experimentation (6), shares a common ground with research describing how interactions between humans could broaden a search (Baumann et al., 2019), but also differs from this prior solution: While interactions between similar agents (i.e., humans) primarily increase organizations' cognitive capacity (Posen et al., 2018), interactions between different (i.e., human and AI) agents also enable the combination of complementary capabilities. These complementarities could make human-AI interaction less prone to joint myopia, which usually limits human interaction's ability to broaden a search (Knudsen & Srikanth, 2014).

Third, combinations of predictive and generative AI enable two mechanisms that allow organizations to reduce their selection bias (Knudsen & Levinthal, 2007). The first mechanism is AI agents' anticipatory quantification (7), which is largely a substitute for humans' heuristic selection, thereby reducing bias toward local solutions. The second mechanism, interactive selection (8), complements humans'



**FIGURE 3**  
**Collective Human Search versus Hybrid Problem-Solving: Outcomes**



*Note:* Figure 3 provides an illustrative representation of different problem-solving types' outcomes. Collective human search, which serves as our baseline, leads mostly to local outcomes (low scope or moderate depth) and some more distant ones when the initial local search fails (illustrated by the spike). In respect of the hybrid problem-solving types, the darker areas indicate where most of the final solutions are expected.

heuristic selection with AI's anticipatory quantification, which counterbalances humans' natural tendencies when selecting.

Our work's second contribution is a descriptive and explanatory model of hybrid problem-solving (see Figure 2). This model predicts that organizations' use of AI-enabled search does not necessarily result in more distant outcomes, but widens the range of outcomes compared to those that current behavioral search models' human-centric perspective predicts (Cyert & March, 1963; Katila & Ahuja, 2002). As shown in Figure 3, a collective human search is most likely to result in local search outcomes (Knudsen & Srikanth, 2014). Conversely, an autonomous search relies on latent representations, latent space exploration, and anticipatory quantification to generate, on average, more distant search outcomes (Figure 3, upper-left quadrant), while a sequential search uses refined representations and a privileged solution search, leading to more local outcomes (lower-right quadrant). Interactive search's shifting representations, interactive experimentation, and interactive selection result in outcomes that, on average, integrate more new knowledge with

existing solution knowledge (upper-right quadrant). Organizations could *occasionally* generate search outcomes corresponding to those predicted for the three hybrid types on the basis of collective human search, such as distant solutions that emerge from humans' creative genius (Kneeland, Schilling, & Aharonson, 2020). In the digital age, however, they could generate these outcomes more *systematically* by deploying hybrid problem-solving.

Our model also suggests that these distinct search outcomes emerge from different combinations of human and AI agents' capabilities in the problem-solving process. Behavioral search theory has generally focused on how humans' specific cognitive capabilities explain search outcomes (March & Simon, 1958; Posen et al., 2018). Our hybrid problem-solving model shifts the focus from humans' cognitive capabilities to the search process, which prior research often treated as a "black box" by linking humans' search capabilities directly to the search outcomes (Posen et al., 2018: 217). Our hybrid problem-solving model is a first step toward exploring specific search processes and their underlying mechanisms that explain the outcomes.

Finally, we contribute to the literature on AI in management (Agrawal et al., 2018; Raisch & Krakowski, 2021) by expanding its scope from the use of predictive AI for routine tasks to generative AI applications for solving novel problems. Generative AI has gained broader attention recently, following the introduction of large language models (LLMs) like ChatGPT (Hao, 2023), even though organizations have been employing other forms of generative AI for several years.<sup>14</sup> Despite these developments, research in management has thus far focused on predictive AI applications in decision-making (Shrestha et al., 2019) and control (Möhlmann, Zalmanson, Henfridsson, & Gregory, 2021). We extend this work by distinguishing between predictive and generative AI, exploring the interrelations between these AI types, and conceptualizing the associated search mechanisms and processes. This extension partially changes prior studies' assumptions. For instance, generating new data has the potential to reduce input data biases, thereby improving prediction. Furthermore, while previous studies on applying AI to routine tasks emphasized accuracy and reliability (Shrestha et al., 2019), nonroutine tasks, such problem-solving, place greater emphasis on novelty and variability (Winter, 2003). As a result, our work lays the foundations for expanding current research on exploitative AI use within an organization's routines (Murray et al., 2021) to more exploratory applications beyond these routines.

### Limitations and Future Research

We set boundaries to limit our propositions and provide the theoretical development focus and depth. Reconsidering these boundaries is beyond our study's scope, but could inform future research. One such boundary refers to our assumptions about sufficient data availability and searching a given space. In practice, organizations may sometimes attempt to use AI in situations where the available data are insufficient—for example, in existing domains where the data are severely biased, or in

new domains where there are few prior solution examples. In these situations, humans could, on the basis of their intuition and contextual understanding, play an important role when auditing AI's inputs, processing, and outputs (Anthony, 2021), or when generating solutions by using abstractions and analogies despite having only limited data (Mitchell, 2019). They could also enable search beyond the current search space by using their creativity to identify new search spaces. For example, humans could combine two databases from distinct fields to create a new search space with more distant solutions (e.g., artists use such "pre-curation" to generate novel AI artworks; see Elgammal, 2019). Future research should therefore explore the role of humans in preparing and enabling the use of AI in situations where the initial data are scarce or biased, or creativity is needed to explore beyond the current search space.

Another boundary is our model's assumption that the different problem-solving types address the same problem. While this is realistic and makes these types comparable, organizations could also select strategically from them. As organizations increase their understanding of the hybrid types' outcomes, they could select those types that are most likely to deliver the desired results for a given problem. However, further research is needed to ascertain this supposition.

A third boundary is the focus on specific problem-solving initiatives. While solving specific problems occurs outside operational routines (Winter, 2003), organizations develop search routines that shape their general approaches to solving problems (Nelson & Winter, 1982: 133; Nigam, Huising, & Golden, 2016). If an organization, for example, develops routines that prioritize an autonomous search, its overall outcomes could exhibit greater novelty than those that organizations primarily using sequential searches produce. Alternatively, organizations could strategically balance their use of different problem-solving types across initiatives. For example, they could balance the use of autonomous search, which is conducive to more distant outcomes, with that of sequential search, which enables more local ones. This strategic use could provide organizations with a new mechanism to balance the dual exploration and exploitation requirements (March, 1991; O'Reilly & Tushman, 2008).

A fourth boundary is our narrow focus on problem-solving. While problem-solving and decision-making are distinct processes (March & Simon, 1958: 160), they could be related. For example, Raisch and Krakowski (2021) suggested that organizations could initially use AI to explore novel problems,

<sup>14</sup> In this paper we have focused on GANs, the most common form of generative AI in practical applications. In contrast to GANs, LLMs do not incorporate a predictor but rely on a pretrained model of language understanding, often referred to as the foundation model. The absence of a built-in predictor renders LLMs less accurate in their predictions. Experts anticipate that combining LLMs with predictors will be necessary to ensure their accuracy and robustness for high-stakes problem-solving contexts, such as drug discovery (Hansen, 2023; Savage, 2023).

but subsequently routinize their decision-making in similar situations once these problems are well-understood and their solutions are known. Future research could investigate how AI applications to solve problems and make decisions affect one another.

A fifth boundary is that we did not explore the consequences beyond the search outcomes. The use of AI-based search alters the work of managers, scientists, and engineers in organizations. Scholars could explore how this affects their job profiles and skill requirements (Krakowski et al., 2023), as well as their work relationships (Kellogg et al., 2020). They could also study the societal implications of AI-based search's use, which has the potential to solve grand challenges, such as finding drugs for neglected diseases, reducing waste and carbon emissions, and designing sustainable products, but could also have negative externalities. For example, a pharmaceutical company's data scientists recently repurposed their employer's drug discovery system. By simply inverting the AI models' parameters—to search for chemical compounds with a *high* toxicity—they identified 40,000 new potential chemical weapons within six hours, including some predicted to be more toxic than any known nerve agent (Urbina, Lentzos, Invernizzi, & Ekins, 2022). This example highlights the need to prevent the misuse of AI-enabled search.

A sixth boundary is that we concentrated on AI technologies currently in use. AI technologies are in a state of constant evolution, and new technologies inevitably emerge. Specifically, it would be intriguing for future research to delve into the emerging utilization of LLMs, such as ChatGPT, in the context of search and problem-solving. Such research could entail comparing and contrasting the application processes and outcomes of these models with those we have described.

Finally, our theoretical ideas need to be empirically explored. Future research could use qualitative methods, such as case studies and ethnographies, which allow for rich insights into longitudinal processes and the organizational contexts in which they occur (Langley, 1999). While these methods are adequate for describing hybrid problem-solving processes in greater depth, field experiments could provide quantitative evidence of their outcomes. Field experiments allow for manipulating the use of different AI implementations in real-life contexts, thereby offering causal insights into outcomes' underlying drivers (Krakowski et al., 2023). For example, a company's marketing teams could use

different hybrid problem-solving types to develop social media content, which allows the comparison of outcomes against those of a control group using collective human search. Alternatively, researchers could rely on archival data on companies' search portfolios, such as the knowledge repositories of pharmaceutical companies' drug discovery activities.

## CONCLUSION

Csaszar and Steinberger (2021) have reminded management scholars that some of their key concepts, like representations and search, originate from foundational AI research. Management scholars, however, have applied these concepts to describe human behavior, largely ignoring their technology heritage. Our study returns to the origins, particularly Simon's (1965) vision of AI in management. Simon completed his pioneering work at a time when the use of AI to solve problems was merely a vision. Today, we build theory from a different vantage point by observing AI-enabled search in organizations. In managerial practice, AI agents are not as almighty as they were in Simon's (1965) vision, since humans and AI agents work jointly on solving problems. These hybrid processes enable rich combinations of human intelligence and AI, which could increase the variety of search outcomes compared to those that emerged in the past. Promising applications suggest that these developments enable organizations to solve problems more effectively, but also generate challenges requiring further management research attention.

## REFERENCES

- Afuah, A., & Tucci, C. L. 2012. Crowdsourcing as a solution to distant search. *Academy of Management Review*, 37: 355–375.
- Agrawal, A., Gans, J., & Goldfarb, A. 2018. *Prediction machines: The simple economics of artificial intelligence*. Boston, MA: Harvard Business School Press.
- Allen, R., & Choudhury, P. 2022. Algorithm-augmented work and domain expertise: The countervailing forces of ability and aversion. *Organization Science*, 33: 149–169.
- Amabile, T. 2020. Creativity, artificial intelligence, and a world of surprises. *Academy of Management Discoveries*, 6: 351–354.
- Anantrasirichai, N., & Bull, D. 2021. Artificial intelligence in the creative industries: A review. *Artificial Intelligence Review*, 55: 1–68.
- Anthony, C. 2021. When knowledge work and analytical technologies collide: The practices and consequences

- of black boxing algorithmic technologies. *Administrative Science Quarterly*, 66: 1173–1212.
- Arthur, R. 2018. Artificial Intelligence empowers designers in IBM, Tommy Hilfiger, and FIT collaboration. *Forbes*.
- Bader, V., & Kaiser, S. 2019. Algorithmic decision-making? The user interface and its role for human involvement in decisions supported by artificial intelligence. *Organization*, 26: 655–672.
- Balasubramanian, N., Ye, Y., & Xu, M. 2022. Substituting human decision-making with machine learning: Implications for organizational learning. *Academy of Management Review*, 47: 448–465.
- Barr, P. S., Stimpert, J. L., & Huff, A. S. 1992. Cognitive change, strategic action, and organizational renewal. *Strategic Management Journal*, 13: 15–36.
- Baumann, O. 2015. Distributed problem solving in modular systems: The benefit of temporary coordination neglect. *Systems Research and Behavioral Science*, 32: 124–136.
- Baumann, O., Schmidt, J., & Stieglitz, N. 2019. Effective search in rugged performance landscapes: A review and outlook. *Journal of Management*, 45: 285–318.
- Bengio, Y. 2012. Deep learning of representations for unsupervised and transfer learning. *JMLR Workshops and Conference Proceedings*, 27: 17–36.
- Bengio, Y., Courville, A., & Vincent, P. 2013. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35: 1798–1828.
- Billinger, S., Benincasa, S., Baumann, O., Kretschmer, T., & Schumacher, T. R. 2023. Learning to search collaboratively: How dyads overcome complexity and misaligned incentives in imperfect modular decompositions. *Industrial and Corporate Change*, 32: 208–233.
- Billinger, S., Srikanth, K., Stieglitz, N., & Schumacher, T. R. 2021. Exploration and exploitation in complex search tasks: How feedback influences whether and where human agents search. *Strategic Management Journal*, 42: 361–385.
- Billinger, S., Stieglitz, N., & Schumacher, T. R. 2014. Search on rugged landscapes: An experimental study. *Organization Science*, 25: 93–108.
- Blohm, I., Antretter, T., Sirén, C., Grichnik, D., & Wincent, J. 2020. It's a people's game, isn't it? A comparison between the investment returns of business angels and machine learning algorithms. *Entrepreneurship Theory and Practice*, 46: 1054–1091.
- Bond, A. H., & Gasser, L. 2014. *Readings in distributed artificial intelligence*. San Matteo, CA: Morgan Kaufman.
- 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.
- Burrell, J. 2016. How the machine “thinks”: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3: 1–12.
- Choudhury, P., Starr, E., & Agarwal, R. 2020. Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal*, 41: 1381–1411.
- Chow, A. R. 2020. There's a wide-open horizon of possibility: Musicians are using AI to create otherwise impossible new songs. *Time*.
- Cortex. 2022. Visit Utah discovers what content skiers love and boosts engagements by 23%. Retrieved from <https://www.meetcortex.com/visit-utah-case-study>.
- 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.
- Csaszar, F. A., & Steinberger, T. 2021. Organizations as artificial intelligences: The use of artificial intelligence analogies in organization theory. *Academy of Management Annals*, 16: 1–37.
- Cyert, R. M., & March, J. G. 1963. *A behavioral theory of the firm*. Englewood Cliffs, NJ: Prentice-Hall.
- Dempsey, M. 2021. Why artificial intelligence is being used to write adverts. *BBC News*.
- Deng, J., Yang, Z., Ojima, I., Samaras, D., & Wang, F. 2022. Artificial intelligence in drug discovery: Applications and techniques. *Brief Bioinform*, 23: bbab430.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. 2015. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology. General*, 144: 114–126.
- Dubin, R. 1978. *Theory building*. New York, NY: Free Press.
- Elgammal, A. 2019. AI is blurring the definition of artist: Advanced algorithms are using machine learning to create art autonomously. *American Scientist*.
- Ethiraj, S. K., & Levinthal, D. 2009. Hoping for A to Z while rewarding only A: Complex organizations and multiple goals. *Organization Science*, 20: 4–21.
- 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.
- Faraj, S., Pachidi, S., & Sayegh, K. 2018. Working and organizing in the age of the learning algorithm. *Information and Organization*, 28: 62–70.
- Fernandes, P., Correia, J., & Machado, P. 2019. Evolutionary latent space exploration and generative adversarial networks. In *Proceedings of the International*

- Conference on the Applications of Evolutionary Computation**, 595–609; Cham, Switzerland: Springer.
- Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Science*, 47: 117–132.
- Fleming, L., & Sorenson, O. 2004. Science as a map in technological search. *Strategic Management Journal*, 25: 909–928.
- Fleming, N. 2018. How artificial intelligence is changing drug discovery. *Nature*, 557: 55–57.
- Fortnow, L. 2013. *The golden ticket: P, NP, and the search for the impossible*. Princeton, NJ: Princeton University Press.
- Furlan, A., Galeazzo, A., & Paggiaro, A. 2019. Organizational and perceived learning in the workplace: A multilevel perspective on employees' problem solving. *Organization Science*, 30: 280–297.
- Gary, M. S., & Wood, R. E. 2011. Mental models, decision rules, and performance heterogeneity. *Strategic Management Journal*, 32: 569–594.
- Gavetti, G., Greve, H. R., Levinthal, D. A., & Ocasio, W. 2012. The behavioral theory of the firm: Assessment and prospects. *Academy of Management Annals*, 6: 1–40.
- Gavetti, G., & Levinthal, D. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45: 113–137.
- Goldberg, D. E. 1989. *Genetic algorithms in search, optimization, and machine learning*. New York, NY: Addison-Wesley Publishing Company.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. 2014. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, vol. 27: 2672–2680.
- Goodfellow, I., Bengio, Y., & Courville, A. 2016. *Deep learning*. Cambridge, MA: MIT Press.
- Gregory, A. 2021. Scientists use AI to create drug regime for rare form of brain cancer in children. *Guardian*.
- Gregory, R. W., Henfridsson, O., Kaganer, E., & Kyriakou, H. 2021. The role of artificial intelligence and data network effects for creating user value. *Academy of Management Review*, 46: 534–551.
- Greve, H. R. 2003. *Organizational learning from performance feedback: A behavioral perspective on innovation and change*. Cambridge, U.K.: Cambridge University Press.
- Hale, C. 2021a. Insilico Medicine begins first human trial of its AI-designed drug for pulmonary fibrosis. *Fierce Biotech*.
- Hale, C. 2021b. Breaking Big Pharma's AI barrier: Insilico Medicine uncovers novel target, new drug for pulmonary fibrosis in 18 months. *Fierce Biotech*. Retrieved from <https://www.fiercebitech.com/medtech/breaking-big-pharma-s-ai-barrier-insilico-medicine-uncovers-novel-target-new-drug-for>.
- Hansen, U. S. 2023. We must perfect predictive models for generative AI to deliver on the AI revolution. *Venture Beat*. Retrieved from <https://venturebeat.com/ai/we-must-perfect-predictive-models-for-ai-to-deliver-on-the-AI-revolution>.
- Hao, K. 2023. What is ChatGPT? What to know about the AI chatbot. *Wall Street Journal*.
- Heath, C., & Staudenmayer, N. 2000. Coordination neglect: How lay theories of organizing complicate coordination in organizations. *Research in Organizational Behavior*, 22: 153–191.
- Holzinger, A. 2016. Interactive machine learning for health informatics: When do we need the human-in-the-loop? *Brain Informatics*, 3: 119–131.
- Jebara, T. 2004. *Machine learning: Discriminative and generative*. Boston, MA: Springer.
- Jeppesen, L. B., & Lakhani, K. R. 2010. Marginality and problem-solving effectiveness in broadcast search. *Organization Science*, 21: 1016–1033.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figumov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Zidek, A., Potapenko, A., Bridgland, A. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596: 583–589.
- Kaplan, S. 2011. Research in cognition and strategy: Reflections on two decades of progress and a look to the future. *Journal of Management Studies*, 48: 665–695.
- Katila, R., & Ahuja, G. 2002. Something old, something new: A longitudinal study of search behavior and new product introduction. *Academy of Management Journal*, 45: 1183–1194.
- Kellogg, K. C., Valentine, M. A., & Christin, A. 2020. Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14: 366–410.
- Kneeland, M. K., Schilling, M. A., & Aharonson, B. S. 2020. Exploring uncharted territory: Knowledge search processes in the origination of outlier innovation. *Organization Science*, 31: 535–557.
- Knudsen, T., & Levinthal, D. A. 2007. Two faces of search: Alternative generation and alternative evaluation. *Organization Science*, 18: 39–54.
- Knudsen, T., & Srikanth, K. 2014. Coordinated exploration: Organizing joint search by multiple specialists to overcome mutual confusion and joint myopia. *Administrative Science Quarterly*, 59: 409–441.
- Krakowski, S., Luger, J., & Raisch, S. 2023. Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*, 44: 1425–1452.

- Krok, A. 2018. GM looks to 3D printing and algorithms to improve future cars. *CNET*.
- Kuchler, H. 2022. Will AI turbocharge the hunt for new drugs? *Financial Times*.
- Langley, A. 1999. Strategies for theorizing from process data. *Academy of Management Review*, 24: 691–710.
- Laursen, K. 2012. Keep searching and you'll find: What do we know about variety creation through firms' search activities for innovation? *Industrial and Corporate Change*, 21: 1181–1220.
- Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. 2022. To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science*, 33: 126–148.
- LeCun, Y., Bengio, Y., & Hinton, G. 2015. Deep learning. *Nature*, 521: 436–444.
- Levinthal, D., & Posen, H. E. 2007. Myopia of selection: Does organizational adaptation limit the efficacy of population selection? *Administrative Science Quarterly*, 52: 586–620.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science*, 43: 934–950.
- Linardatos, P., Papastefanopoulos, V., & Kotsiantis, S. 2020. Explainable AI: A review of machine learning interpretability methods. *Entropy*, 23: 1–45.
- Lindebaum, D., Vesa, M., & den Hond, F. 2020. Insights from “The machine stops” to better understand rational assumptions in algorithmic decision making and its implications for organizations. *Academy of Management Review*, 45: 247–263.
- Luo, X., Qin, M. S., Fang, Z., & Qu, Z. 2021. Artificial intelligence coaches for sales agents: Caveats and solutions. *Journal of Marketing*, 85: 14–32.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science*, 2: 71–87.
- March, J. S., & Simon, H. A. 1958. *Organizations*. New York, NY: Wiley.
- Metcalf, L., Askay, D. A., & Rosenberg, L. B. 2019. Keeping humans in the loop: pooling knowledge through artificial swarm intelligence to improve business decision making. *California Management Review*, 61: 84–109.
- Metz, C. 2020. How A.I. steered doctors toward a possible coronavirus treatment. *New York Times*.
- Mitchell, M. 2019. *Artificial intelligence: A guide for thinking humans*. London, U.K.: Penguin.
- Möhlmann, M., Zalmanson, L., Henfridsson, O., & Gregory, R. W. 2021. Algorithmic management of work on online labor platforms: When matching meets control. *Management Information Systems Quarterly*, 45: 1999–2022.
- Murray, A., Rhymer, J., & Sirmon, D. G. 2021. Humans and technology: Forms of conjoint agency in organizations. *Academy of Management Review*, 46: 552–571.
- Nelson, R., & Winter, S. 1982. *An evolutionary theory of economic change*. Cambridge, MA: Harvard University Press.
- Newell, A., Shaw, J. C., & Simon, H. A. 1959. Report on a general problem-solving program. *Proceedings of the International Conference on Information Processing*, 256–264: Paris, France: UNESCO.
- Newman, D. T., Fast, N. J., & Harmon, D. J. 2020. When eliminating bias isn't fair: Algorithmic reductionism and procedural justice in human resource decisions. *Organizational Behavior and Human Decision Processes*, 160: 149–167.
- Ng, A., & Jordan, M. 2001. On discriminative vs. generative classifiers: A comparison of logistic regression and naïve Bayes. In *Proceedings of the 14th International Conference on Neural Information Processing Systems*, 841–848: Cambridge, MA: MIT Press.
- Nigam, A., Huising, R., & Golden, B. 2016. Explaining the selection of routines for change during organizational search. *Administrative Science Quarterly*, 61: 551–583.
- Oberhaus, D. 2020. NASA's new moon-bound space suits will get a boost from AI. *Wired*.
- Oh, S., Jung, Y., Kim, S., Lee, I., & Kang, N. 2019. Deep generative design: Integration of topology optimization and generative models. *Journal of Mechanical Design*, 141: 11–33.
- O'Reilly, C. A., III, & Tushman, M. L. 2008. Ambidexterity as a dynamic capability: Resolving the innovator's dilemma. *Research in Organizational Behavior*, 28: 185–206.
- Orlikowski, W. J., & Scott, S. V. 2014. What happens when evaluation goes online? Exploring apparatuses of valuation in the travel sector. *Organization Science*, 25: 868–891.
- Ott, T. E., Eisenhardt, K. M., & Bingham, C. B. 2017. Strategy formation in entrepreneurial settings: Past insights and future directions. *Strategic Entrepreneurship Journal*, 11: 306–325.
- Piezunka, H., & Dahlander, L. 2015. Distant search, narrow attention: How crowding alters organizations' filtering of suggestions in crowdsourcing. *Academy of Management Journal*, 58: 856–880.
- Posen, H. E., Keil, T., Kim, S., & Meissner, F. D. 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.

- Posen, H. E., & Martignoni, D. 2018. Revisiting the imitation assumption: Why imitation may increase, rather than decrease, performance heterogeneity. *Strategic Management Journal*, 39: 1350–1369.
- Prosser, D. 2021. How Pencil aims to automate the secret sauce of advertising. *Forbes*.
- Puranam, P. 2021. Human-AI collaborative decision-making as an organization design problem. *Journal of Organization Design*, 10: 75–80.
- 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.
- Raisch, S., & Krakowski, S. 2021. Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review*, 46: 192–210.
- Rivkin, J. W., & Siggelkow, N. 2003. Balancing search and stability: Interdependencies among elements of organizational design. *Management Science*, 49: 290–311.
- Russell, S., & Norvig, P. 2020. *Artificial intelligence: A modern approach*. Boston, MA: Pearson.
- Savage, N. 2023. Drug discovery companies are customizing ChatGPT: Here is how. *Nature Biotechnology*, 41: 585–586.
- Schilling, M. A., & Green, E. 2011. Recombinant search and breakthrough idea generation: An analysis of high impact papers in the social sciences. *Research Policy*, 40: 1321–1331.
- Schwab, K. 2019. This is the first commercial chair made using generative design. *Fast Company*.
- Senoner, J., Netland, T., & Feuerriegel, S. 2022. Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing. *Management Science*, 68: 5704–5723.
- Shorten, C., & Khoshgoftaar, T. M. 2019. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6: 1–48.
- Shrestha, Y. R., Ben-Menahem, S. M., & Von Krogh, G. 2019. Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61: 66–83.
- Shrestha, Y. R., Krishna, V., & Von Krogh, G. 2021. Augmenting organizational decision-making with deep learning algorithms: Principles, promises, and challenges. *Journal of Business Research*, 123: 588–603.
- Simon, H. A. 1957. *Models of man: Social and rational*. New York, NY: Wiley.
- Simon, H. A. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106: 467–482.
- Simon, H. A. 1965. *The shape of automation for men and management*. New York, NY: Harper & Row.
- Simon, H. A., Dantzig, G. B., Hogarth, R., Plott, C. R., Raiffa, H., Schelling, T. C., Shepsie, K. A., Thaler, R., Tversky, A., & Winter, S. 1987. Decision making and problem solving. *Interfaces*, 17: 11–31.
- Simone, S. 2021. Pencil offers new AI-driven ad creation service to help small businesses. *KM World*.
- Stuart, T. E., & Podolny, J. M. 1996. Local search and the evolution of technological capabilities. *Strategic Management Journal*, 17: 21–38.
- Tanaka, F. H. K. D. S., & Aranha, C. 2019. Data augmentation using GANs. *arXiv*: 1904.09135.
- Toews, R. 2021. AlphaFold is the most important achievement in AI – Ever. *Forbes*.
- Tripsas, M., & Gavetti, G. 2000. Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management Journal*, 21: 1147–1161.
- Uotila, J., Keil, T., & Maula, M. 2017. Supply-side network effects and the development of information technology standards. *Management Information Systems Quarterly*, 41: 1207–1226.
- Urbina, F., Lentzos, F., Invernizzi, C., & Ekins, S. 2022. Dual use of artificial-intelligence-powered drug discovery. *Nature Machine Intelligence*, 4: 189–191.
- Vinoski, J. 2019. Generative design opens new paths to a sustainable future. *Forbes*.
- Von Krogh, G. 2018. Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Academy of Management Discoveries*, 4: 404–409.
- Walsh, T. 2017. The singularity may never be near. *AI Magazine*, 38: 58–62.
- Winter, S. G. 2003. Understanding dynamic capabilities. *Strategic Management Journal*, 24: 991–995.
- You, S., Yang, C. L., & Li, X. 2022. Algorithmic versus human advice: Does presenting prediction performance matter for algorithm appreciation? *Journal of Management Information Systems*, 39: 336–365.
- Zahra, S. A., & George, G. 2002. Absorptive capacity: A review, reconceptualization, and extension. *Academy of Management Review*, 27: 185–203.



**Sebastian Raisch** (sebastian.raisch@unige.ch) is a professor of strategy at the Geneva School of Economics and Management, University of Geneva. He received his PhD from the University of Geneva and his habilitation from

the University of St. Gallen. His research interests include artificial intelligence in management, organizational ambidexterity, and organizational paradox.

**Kateryna Fomina** (kateryna.fomina@unige.ch) is a doctoral student at the Geneva School of Economics and

Management, University of Geneva. Her research explores the collaboration between humans and artificial intelligence in organizations.

