



8

A Brief History of the Relationship Between Expertise and Artificial Intelligence

Jan Maarten Schraagen and Jurriaan van Diggelen

This chapter explores the role artificial intelligence (AI) plays in human expertise, for instance by either enhancing, changing, or degrading it. We also address how expertise can play a role in moderating, advancing, using, collaborating with, or exploiting AI. We should make clear that we will neither set up a simple dichotomy between experts and AI, nor will we investigate claims of people being surpassed in expertise by artificial general intelligence (AGI), or people becoming unemployable due to AI developments. However, we do not deny the partial validity of some of these claims. Rather, we view experts and AI systems as ‘joint cognitive systems’ that form a unit (Woods & Hollnagel, 2006). There are numerous ways for humans, and experts in particular, to jointly collaborate with AI systems, and we discuss the empirical evidence for particular patterns of collaboration. Moving beyond a ‘joint cognitive systems’ approach, we

J. M. Schraagen (✉) • J. van Diggelen
TNO (The Netherlands Organisation for Applied Scientific Research),
Soesterberg, The Netherlands
e-mail: jan_maarten.schraagen@tno.nl; jurriaan.vandiggelen@tno.nl

also discuss more recent ways in which AI has manifested itself as a networked and distributed phenomenon and has shown itself to either enhance or degrade human expertise. To achieve this, we first present a brief history of AI and expertise studies. Next, we provide examples of empirical research on experts working together with intelligent systems and emphasize the patterns that emerge from that research to shed light on the role of AI in expertise. Subsequently, we discuss a case study in radiology that illustrates how human experts and AI approach this topic. Finally, we conclude and provide some recommendations for future research.

The concepts of expertise, intelligence, and artificial intelligence are used frequently in this chapter. The distinction between expertise and intelligence is one between domain-specific and domain-generic knowledge (Vergne, 2017). Typically, expertise is defined in terms of “reliably superior performance on representative tasks” (Ericsson, 2006, p. 13), although this definition is arguably more applicable to tasks that can be measured, standardized, or simulated easily (e.g., chess, music, typing, or playing tennis) rather than complex cognitive work where performance measurement is difficult or impossible (Ward et al., 2020). Intelligence, in contrast, may be defined as “a very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly, and learn from experience” (Gottfredson, 1997, p. 13). Evidence shows that intelligence is a reasonably good predictor of performance early in learning, but does not predict asymptotic levels of learning very well (Hunt, 2006). In a recent review, Hambrick, Burgoyne, and Oswald (2020) concluded that the evidence for the role of general cognitive ability in expertise is inconclusive and in the majority of studies the evidence was in fact absent. On the other hand, cognitive ability did play a role in job performance well beyond the initial training. The difference between expertise and job performance studies is that the former typically studies consistent mappings between stimulus and response (as in the routine execution of psychomotor responses or the recognition of typical patterns of stimuli), whereas the latter involves acquiring new knowledge and skill, dealing with varied mappings between stimulus and response or the need to

develop mental models of a situation. Thus, general cognitive ability (of which intelligence is one construct) plays a role whenever the environment presents us with new or complex situations. Whenever the environment presents us with well-known, standardized situations, we draw upon domain-specific knowledge and call it ‘expertise’.

The European Union High-Level Expert Group on Artificial Intelligence recently provided an updated definition of AI, which we use in this chapter:

Artificial intelligence (AI) systems are software (and possibly also hardware) systems designed by humans that, given a complex goal, act in the physical or digital dimension by perceiving their environment through data acquisition, interpreting the collected structured or unstructured data, reasoning on the knowledge, or processing the information, derived from this data and deciding the best action(s) to take to achieve the given goal. AI systems can either use symbolic rules or learn a numeric model, and they can also adapt their behavior by analyzing how the environment is affected by their previous actions. (European Commission, 2019, p. 6)

AI systems achieve intelligent, that is, rational behavior, by choosing the best action to take in order to achieve a certain goal. Current AI systems can be characterized as narrow AI systems which perform one or a few specific tasks and cannot deal well with any new or abnormal situation. These systems resemble our definition of expertise as “reliably superior performance on representative tasks” (which is not to say that narrow AI systems should be equated with human experts, as the latter also possess general cognitive ability that the narrow AI systems by definition do not).

Taking these definitions into account, discussing the role of AI in expertise can mean a number of things. Given that currently deployed AI systems are examples of narrow AI, the issue becomes one of how human experts, within their domain of expertise, work together with systems that can perform one or a few specific tasks within that domain of expertise. In other words, experts work with AI as symbiotic partners to exploit what each party does best (Daugherty & Wilson, 2018).

History, Current Status, and Prospects of Artificial Intelligence

The history of AI is often divided into multiple phases that characterize the field as defined by a particular research interest, or technological success. In this section, we briefly discuss these phases through the lens of understanding the role of AI in expertise. As a guideline, we follow the phases as distinguished by Nilsson (2009) and shown in Table 8.1. We finish the section with a phase describing our expectation for the future.

Early Days (1956–1974)

The field of artificial intelligence was founded during the legendary Dartmouth workshop in 1956. In the years that followed, the workshop participants (among others) developed many of the core techniques and ideas in AI that would continue to exist today. The first important idea was that knowledge could be represented symbolically (at the time referred to as a semantic network by Quillian (1963)), and logic could be used to reason over it. Another important idea was that knowledge could also be represented using a connectionist approach in an artificial neural network (at the time referred to as perceptron), being loosely inspired by the working of the human brain. The wide range of possibilities and various successful early prototypes, such as chess computers, and programs processing natural language led to high expectations of this new emerging field. Prominent researchers such as Herbert Simon and Marvin

Table 8.1 Phases of artificial intelligence (AI)

Years	Characterization of AI
1956–1974	Early days of symbolic reasoning
1974–1980	First AI winter
1980–1987	Expert systems
1987–1993	Second AI winter
1993–2011	Multi-agent systems and semantic web
2011–present	Big data and deep learning
Future paradigm of AI	Hybrid AI

Minsky predicted that AI would surpass human experts on selected tasks within a few decades.

However, progress was hampered by a number of problems. The first problem was the lack of computing power in early computers. The second problem was the burden of manual work required to engineer all the facts and rules required for intelligent reasoning. It gradually became apparent that general search strategies (so-called weak methods) were insufficient for attaining high levels of performance, and that these strategies needed to be complemented with a lot of domain knowledge. The third problem was that AI models turned out to be *brittle*, meaning that they only performed well on the limited scope they were designed for. The latter two problems were conceived as part of the research process: just as in other successful sciences like physics, basic principles should first be investigated using simplified models. Researchers focused on *micro worlds* (Minsky & Papert, 1972), which would be narrow at first, but could later be generalized to more realistic settings. This generalizability turned out to be problematic, hampering practical applications.

First AI Winter (1974–1980)

These problems, coupled with the unrealistically high expectations, led to what is generally called the first AI winter. Research funding was cut, and the general expectations of AI were dramatically lowered. Researchers came to realize that the problem of modeling intelligence in a computer was to be much harder than they initially thought.

Expert Systems (1980–1987)

Following the realization that weak methods were insufficient for realizing high levels of performance, researchers turned to ways of incorporating large amounts of domain knowledge into systems. These systems were called expert systems, as they were assumed to encapsulate the knowledge of experts in a particular domain. Expert systems were building on early insights in symbolic knowledge representation. Knowledge was

represented using production rules (usually handcrafted by human experts), and a reasoning engine was applied to derive consequences given a set of facts. Popular applications were the medical domain (e.g., Mycin) and law. The goal was to “*incorporate the knowledge and expertise in computer programs, making the knowledge and expertise easily replicated, readily distributed, and essentially immortal*” (Davis, 1984, p. 18, our emphasis). Just as in the early days of AI, expectations were high (Bobrow, 1984).

Besides the progress in expert systems, significant advances were also made in the connectionist approach to AI due to the discovery of the multi-layer perceptron that solved one of the fundamental problems of the old perceptron model from the 1960s. However, these developments did not create as much enthusiasm as expert systems, and it was unclear how the two approaches could be combined. Additionally, the problems with expert systems were essentially the same as in the early days: brittleness, and burden of manual work. The main strategy to counter these threats was to limit the application to a narrowly defined topic, avoiding the need to model common-sense knowledge in the system. Another strategy was to try to enable end users to model the expert system rules. Nevertheless, expert systems did not live up to their expectations, and rarely made it out of the lab to real life usage (Leith, 2016).

Second AI Winter (1987–1993)

Similar to the first AI winter, the inability to live up to the high expectations caused a second AI winter. This led many researchers to look for a different paradigm. Some researchers argued for an entirely different approach, referring to the symbolic approach to AI as GOF AI (Good Old-Fashioned AI), which was perceived as fundamentally flawed (Brooks, 1990). Furthermore, the term expert system was replaced by *decision support system* to reflect a ‘downscaled’ ambition where the computer serves as a helper of a human expert instead of being an expert itself.

Multi-agent Systems and the Semantic Web (1993–2011)

Renewed hope in artificial intelligence was raised by a new technology that would fundamentally transform computer science: the internet. One development was multi-agent systems (MAS), which is a paradigm for distributed artificial intelligence. A MAS comprises multiple active AI-entities and lacks a single point of control and can therefore be considered as more robust (potentially overcoming the *brittleness* problem). Furthermore, it allows multiple developers to work on a system with little or no coordination. This was believed to be a potential solution for relieving the burden of work. Another new development was the semantic web, which was viewed as a next step in the evolution of symbolic knowledge representation. The novelty was that it was distributed. Ontologies serve as formal specifications of the conceptualizations that are shared between the knowledge sources (Gruber, 1993). Unfortunately, both MAS and the semantic web did not live up to their high expectations, and few practical applications resulted from it.

Big Data and Deep Learning (2011–Present)

The difficulties of MAS and the semantic web did not result in another AI winter. Large amounts of data (also known as ‘big data’) were created as a result of increased computer memory, sensor technology, and (again) the internet. Big data turned out to be a missing ingredient required to make the connectionist approach work. The large availability of data and computing power made it possible to develop deep neural networks (DNN) with up to one hundred million parameters that automatically optimize using machine learning techniques (many of which had already been discovered decades ago). Deep learning turned out to be very successful, leading to unprecedented outcomes such as superhuman performance on image classification tasks, game-playing such as the board game Go, and major breakthroughs in voice recognition and automatic language translation among many others. For the first time in history, AI

became a huge commercial success, giving rise to billion-dollar industries in highly automated driving and data-analytics.

Not surprisingly, these successes revived speculations about the glorious future of AI, including the possible development of artificial general intelligence (AGI), and super intelligence (Bostrom, 2017). Many people believed that deep learning had finally solved the problem of brittleness and manual engineering, thus making all previous approaches in AI obsolete. With respect to the problems of brittleness and burden of manual work, there has certainly been progress. Advocates of end-to-end DNNs point out that feature extraction (e.g., extracting phonemes in audio) is no longer required. The raw features (e.g., the waveform itself) should be directly fed into the DNN, which should be trained to produce the output in one go. This bypasses the manual engineering of domain-specific feature extraction algorithms. Furthermore, it enhances performance, hence reducing brittleness. However, there are two main problems with this approach, which indicates a fundamental shortcoming of end-to-end deep learning.

First, deep learning requires a lot of data. For an image classifier, requiring one million training examples is common. The problem is that these images must be accompanied by a label. A label could, for example, state that a certain image qualifies as ‘a cat’ and another as ‘a dog’. Because deep learning is a supervised learning algorithm, it requires these labels to learn. To obtain a label, a dataset usually requires humans to point out the area and indicate which type of object resides there. Whereas manually engineering a dataset for highly automated cars may be considered worth the effort, for more rare and specialized applications this burden of manual labeling work is often too large or simply not feasible.

A second problem with end-to-end DNNs is that they are no longer understandable by humans. The network cannot explain why it has reached a certain conclusion, which is problematic when humans have to judge the trustworthiness of an AI algorithm’s outcome. Although much research is currently performed on explainable AI (Gunning & Aha, 2019), this research is still in its infancy and most likely requires more than a DNN to be solvable. Performing calculations with tens of millions of parameters, the functioning of a deep learning network is inherently

incomprehensible to humans. This can lead to unexpected behaviors and errors. For example, researchers discovered that small perturbations in the input image (invisible to the human eye) could easily fool a neural image classifier (e.g., confusing a whale with a turtle) (Moosavi-Dezfooli, Fawzi, & Frossard, 2016). The network turned out to be brittle after all, but in a way that is totally unimaginable for humans and that could not be explained by the AI either.

Hybrid AI (The Future Paradigm of AI)

Whereas deep learning undoubtedly has proven its usefulness in pattern recognition tasks, many believe that the approach is not extendable to more complex tasks (Marcus & Davis, 2019). For example, consider an AI algorithm that could predict whether a business strategy will be successful or not. Imagine an end-to-end DNN that takes a description of a strategy and situation as input, and produces an output that labels the strategy as 'good' or 'bad'. As attractive as such a solution may seem, the data to train such a network are simply not available in the right format and quantity. Furthermore, the output will probably never be that black and white, requiring the algorithm to explain its advice, something which DNNs are inherently poor at.

While no one can predict the future, we believe that a future AI era will go beyond deep learning (Peeters et al., 2020). In fact, its contours are already beginning to take shape. In this era, AI will evolve into a hybrid of multiple connectionist AI techniques, symbolic approaches, and humans. By merging symbolic and connectionist approaches (van Harmelen & Teije, 2019), a hybrid AI system can be developed, which combines human-understandability and high-level reasoning with pattern recognition capabilities. Furthermore, humans will also become an essential part of the system fulfilling essential roles as bearers of responsibility, handling unexpected situations that the AI is incapable of deriving, and discovering causal relationships that are not discoverable by observing data alone (Pearl & Mackenzie, 2018).

History of Expertise Studies

In the 1950s and 1960s, research on expertise, particularly in the United States, was relatively scarce. Woodworth and Schlosberg's (1954) *Experimental Psychology* does not mention the topic at all. One of the few exceptions was the work on chess expertise by the Dutch psychologist Adriaan de Groot (1946, 1965). De Groot collected think-aloud protocols of chess players of varying expertise between 1938 and 1943. Although many at the time thought there would be large differences in the number of moves considered or the depth of search between grandmasters and amateurs, de Groot found no evidence for such differences. However, he did find differences in the speed with which complex board positions could be stored in memory and remembered correctly after being presented for only five seconds. Chess masters could correctly reconstruct positions of more than 20 pieces after just five seconds of study, whereas the amateurs could reconstruct only four or five pieces. Apparently, the chess masters were able to recognize meaningful patterns on the board, later called 'chunks', indicating that domain-specific chess knowledge was the determining factor in the observed difference between experts and beginners.

The work by de Groot turned out to be highly influential and foundational once it was translated into English in 1965. Around this time, research in AI reached a dead end in that it had failed to construct computer programs that could outperform humans (Feigenbaum, 1989; Glaser & Chi, 1988). The weak search methods implemented in these programs employed heuristics to prune exhaustive search trees, but to no avail. Although heuristics are knowledge, they are a form of general knowledge. Looking at this state of affairs with de Groot's findings in mind, researchers became aware of the importance of domain-specific knowledge in expertise. Chess masters don't differ from amateurs because of their efficient wielding of general search heuristics, but because of their large storage of knowledge of chess patterns and associated moves. Simon and Gilmartin (1973) estimated that masters have acquired on the order of 50,000 different chess patterns, that they can quickly recognize such patterns on a chessboard and that this ability is what underlies their superior performance in chess.

The ‘classic expertise approach’ (for an overview see Gobet, 2020) started with the originating work by Chase and Simon (1973) on chess at Carnegie-Mellon University in the early 1970s. This approach is characterized by detailed analyses of problem-solving processes by a relatively small number of participants, emphasis on content, and use of computer programs to express theories. Chase and Simon also introduced a variation on de Groot’s memory task, basically serving as a control condition: apart from presenting actual board configurations, participants were also given random board configurations. In the latter case, no differences were observed between experts and beginners (Chase & Simon, 1973). This showed that the results obtained with actual board configurations were not due to superior visual memory for isolated pieces, but rather depended critically upon the ‘meaning’ of the constellations of pieces (‘chunks’). This research spawned a flurry of experimental papers in the late 1970s and early 1980s that would be summarized by Anderson (1981) and Chi, Glaser, and Farr (1988). Not only was the skill effect in the memory recall task replicated in several domains, but it was also found that experts see and represent a problem in their domain at a deeper (more principled) level than novices (Chi, Feltovich, & Glaser, 1981).

In 1991, Holyoak asserted that “[t]heories of expertise have now passed through two generations” (p. 301). The first generation viewed expertise as essentially a problem-solving activity that employed general heuristic search methods (akin to the ‘weak methods’ discussed previously) to a broad range of domains. However, in the 1970s and early 1980s, it became clear that expertise depended crucially on extensive domain knowledge and was therefore limited in scope and did not transfer across domains (for an overview see Feltovich, Prietula, & Anders Ericsson, 2006). Interestingly, the field of AI had gone through a similar major shift in focus in the 1966–1976 period, essentially moving from a search paradigm to a knowledge-based one (Goldstein & Papert, 1977), culminating in the heyday of highly domain-specific expert systems (Feigenbaum, McCorduck, & Nii, 1988). It seemed clear from all of this research that “knowledge is power” (Feigenbaum, 1989), which captured the essence of the second generation of theories of expertise.

Yet, in 1991, Holyoak listed numerous empirical findings that were at odds with the second generation of expertise theories. He found that

experts were much more flexible than previously thought and summarized his findings by stating that “[i]n general, an expert will have succeeded in adapting to the inherent constraints of the task” (Holyoak, 1991, p. 309). In other words, rather than reliably attaining specific goals within a specific domain (the second-generation definition of ‘routine expertise’), expertise should be viewed as the ability to make an appropriate response to a situation that contains a degree of unpredictability. The latter definition of expertise was first advanced by Hatano and Inagaki (1986) and was called adaptive expertise. Holyoak (1991) went on to outline a connectionist view of expertise. However, he did not convincingly demonstrate that a symbolic connectionist approach could explain the empirical findings that were at odds with the second-generation theories of expertise and this approach to expertise was not taken up widely (it may have been before its time). In fact, the classic expertise approach has remained one of the dominant approaches to expertise (Gobet, 2020) and has been extended to expert decision making in real-world situations in the field of Naturalistic Decision Making (see Schraagen, 2018, for how this field relates to the theoretical foundations laid by the classic approach to expertise). In the field of Human Resource Development, the classic expertise approach, with its focus on knowledge, experience, and problem solving, has been extended with subjective characteristics that are perceived by someone else as an indication of an expert’s knowledge, abilities, or skills, for instance, being motivated, self-confident, or having high interpersonal skills (Germain, 2006; Germain & Tejada, 2012; Grenier & Germain, 2014).

Currently, there is no single overarching and commonly accepted definition of expertise. In the recent *Oxford Handbook of Expertise*, Ward et al. (2020) distinguish many communities of practice that all use the word expertise in different ways. Apart from the classic expertise approach, the Cognitive Systems Engineering community of practice offers perhaps the most distinctive alternative. It does not view expertise as an individual phenomenon or a particular stage of information processing, as the classic expertise approach, but rather as a coupling between an expert with a problem ecology through a representation. In this view, expertise is a matter of sensitivity to environmental constraints and opportunities.

The pendulum on the generality-specificity dimension has therefore swung back to some extent, and many researchers now view expertise as “skilled adaptation to complexity and novelty” (Ward, Gore, Hutton, Conway, & Hoffman, 2018), therefore stressing generality somewhat more than specificity. Research has confirmed the importance of conscious, analytical reasoning (as an instance of skilled adaptation or flexibility) in experts, but only when confronted with complex, atypical problems (Mamede et al., 2010; Moxley, Anders Ericsson, Charness, & Krampe, 2012). When having to solve simple problems, experts use a recognition strategy, as predicted by the classic expertise approach, and the first option considered is usually the best (e.g., Johnson & Raab, 2003; Klein, Wolf, Militello, & Zsombok, 1995). The importance of a flexible and adaptive skill capacity (e.g., flexible sensemaking and flexible action execution) will only increase as the societal and human-technological challenges ahead of us proliferate.

Interestingly, whereas flexibility and adaptation are prominent concepts in current conceptualizations of expertise, current conceptualizations of AI still focus on attaining specific goals within a specific domain. Most AI systems of note have so far achieved world-class performance in specific domains such as the competitive games of chess (Campbell, Hoane Jr., & Hsu, 2002), Go (Silver et al., 2017), Jeopardy (Chen, Elenee Argentinis, & Weber, 2016; Ferrucci, 2012), and Poker (Brown & Sandholm, 2018). Nevertheless, they are still far away from what is sometimes referred to as Artificial General Intelligence (AGI), meaning that it can perform any intellectual task that a human can. Having discussed the history of both AI and expertise studies, we will now turn to studies on ‘joint cognitive systems’, in which experts and intelligent systems are viewed as pairs that work together to achieve particular goals.

Empirical Research on Joint Cognitive Systems

An early example of an empirical study on the coupling of human intelligence and machine power is the study by Roth, Bennett, and Woods (1987) on technicians diagnosing faults with the aid of an expert system. The expert system was developed according to what the authors refer to

as a prosthesis paradigm, which may be contrasted with a cognitive instrument paradigm. In the cognitive tool as a prosthesis paradigm, “[t]he machine expert guides all problem solving activities dictating what observations and actions the user is to take to solve the problem” (Roth et al., 1987, p. 480). The expert system is considered a prosthesis in the sense that it presumably compensates for human deficiencies in generating hypotheses and the human is relegated to the role of passive data gatherer and action implementer in order to serve the machine’s needs.

The study showed that those technicians who were actively involved in the troubleshooting process not only achieved faster and better solutions, but also coped better with unanticipated variability, monitored the machine’s behavior, recognized unproductive paths, and redirected the machine to more productive paths. Technicians who passively followed the machine’s instructions dwelled on unproductive paths and reached dead-ends more often. It turned out that one of the six problems presented to the technicians was unsolvable due to a bug in the expert system’s knowledge base. Substantive interventions by the knowledge engineer were also required to point out input errors or redirect the diagnosis. Technicians varied widely in how they approached the problem and substantial deviations from the canonical path arose even when the problem was solved correctly. It turned out that technicians needed knowledge of the structure and function of the device in order to follow underspecified instructions by the expert system, to infer machine intentions, to resolve impasses, and to recover from errors that led the expert system off-track (once off-track, it could not recover by itself and needed human help to be directed back to a more productive path). In brief, the expert system was not observable, predictable, and directable by the human expert.

The machine-as-prosthesis paradigm results in typical breakdowns in performance whenever humans are assigned the passive role of following instructions. Alternatively, cognitive tools can also be viewed as instruments that support effective performance in any environment. This instrumental view of tools is very much in alignment with the view of expertise as skilled adaptation to complexity and novelty. Tools as instruments should enhance a human problem solver’s adaptability to the unanticipated variability that inevitably arises in the pursuit of domain

goals. The problem solver is in charge, the AI tool functions more as a staff member providing knowledge resources.

This example of a joint cognitive system focused on a single human and a single system, even though it became clear during this particular research project that the scope had to be extended to include the knowledge engineer and two observers who could help and guide the technician where necessary. Later research on joint cognitive systems extended to multiple experts cooperating with multiple intelligent systems. One typical domain would be automation in the airplane cockpit, where the cockpit crew needs to cooperate with numerous automated systems. Not all these systems qualify as artificial intelligence, as some of them hardly 'interpret information' or 'reason based on knowledge', but that is beside the point here. The point that we want to make, and that has been stated repeatedly by the field of Cognitive Systems Engineering (e.g., Woods, Dekker, Cook, Johannesen, & Sarter, 2010; Woods & Hollnagel, 2006), is that a clumsy use of technology is about miscoordination between the human and machine portion of a single ensemble (Christoffersen & Woods, 2002). Automation and people have to coordinate as a joint system, a single team (Klein, Woods, Bradshaw, Hoffman, & Feltovich, 2004). Breakdown in this team's coordination is an important path toward disaster, as can be seen vividly in the Air France Flight 447 disaster (2009) or the Lion Air (2018) and Ethiopian Air (2019) crashes involving the Boeing 737 Max MCAS system.

In essence, what happens with many (cockpit) automation projects is that systems are designed to operate in a multitude of modes, and mode changes are not always communicated clearly to operators. Mode errors occur when an operator executes an intention that is appropriate for one mode, when in fact the system is in a different mode. For instance, when a pilot enters the correct digits for a planned descent (e.g., '33', intending to mean an angle of descent of 3.3 degrees), this may be interpreted by the automation (being in a different descend mode than the pilot thinks) as a rate of descent of 3300 feet per minute. This particular mode error occurred with Air Inter Flight 148 in 1992 near Strasbourg, France, killing 87 of the 96 people on board. On a more day-to-day level, cruise control systems in cars provide opportunities for mode errors as well. For instance, one may manually override the speed set by the cruise control

by pressing the gas pedal for a while, then forgetting about the cruise control being engaged, only to be reminded of it when releasing the gas pedal and letting the car gradually slow down. If one's intention was to slow down to zero mile per hour, the cruise control would suddenly kick in at the set speed, and one will experience an 'automation surprise' (Sarter & Woods, 1995), much like pilots in a cockpit. Drivers may also believe that the cruise control is engaged, when in fact it is only on. How the various modes are communicated to the driver is highly dependent on the particular cruise control interface, and different car manufacturers have different ways of resolving this issue.

Mode errors are only one example of where automation has not lived up to its promise. Other examples are clumsy automation (Wiener, 1989) where automation creates new coordination demands precisely at the very time when practitioners are most in need of true assistance, overreliance on technology (Billings, 1991) where operators rely on systems when in fact those systems cannot cope, as they are outside their competence envelope, and deskilling (Bainbridge, 1983) where operators gradually lose manual skills as they increasingly depend upon automation. Underlying these problems with automation are several misconceptions regarding the way tasks are to be distributed among people and technology (Bradshaw, Hoffman, Johnson, & Woods, 2013):

1. Compensation: machines have strong points that compensate for weak points of humans;
2. Substitution: tasks can be automated without consequences; hence human tasks can be replaced with machine tasks;
3. Automation: automation is autonomous;
4. Allocation: tasks can be neatly divided into parts and assigned to either a human or a machine (not both at the same time); and
5. Workload and productivity: more automation leads to fewer people, hence fewer errors, hence lower costs, but with higher productivity.

Many of the current discussions around AI can be framed as novel instantiations of the same discussions on automation: if one replaces 'machine' or 'automation' with 'AI' in the misconceptions above, one would find themselves in the same position as cognitive engineers in the

1980s and 1990s. Many of the lessons learned then with automation still apply in the case of AI, even though the empirical evidence is still unavailable. The following arguments may be advanced in response to some of the misconceptions:

1. Compensation: machines/AI are good at certain things and people are good at certain things, but that does not change the fundamental interdependence between the two. Team play with people and AI is critical to success. No matter how much information the AI processes, humans must trust the conclusions because they are ultimately responsible. Therefore, AI needs to explain itself.
2. Substitution: practice is transformed by automation and the roles of people change. This may not always be obvious from an outsider's perspective, due to the Law of Fluency that states that 'well'-adapted work occurs with a facility that belies the difficulty of the demands resolved and the dilemmas balanced (Woods & Hollnagel, 2006). In other words, when an outsider studies work that seems to be well-adapted, what remains hidden from view are the numerous ways in which humans have coped with complexity and the various trade-offs they had to make. As the constraints adapted to are hidden from view, the work may actually not be so 'well'-adapted. Humans will adapt to changes in the tasks as a result of automation, but that adaptation comes at a price, for instance deskilling, increased monitoring, or increased coordination. These vulnerabilities will become apparent when situational demands increase, and surprise events occur.
3. Automation: Machines are self-sufficient only up to a certain extent and only in particular circumstances. Surprise is continuous and ever-present. There is always the need to close the gap between the demonstration and the real thing (Woods, 2016). This requires new methods to assess brittleness, for instance the turnaround test—how much work does it take to get a system ready to handle the next mission/case/environment, when the next is not a simple parametric variation of the previous demonstration (Woods, 2016)? As a second rebuttal, it has recently been claimed that “no AI is an island” (Johnson & Vera, 2019). According to Johnson and Vera (2019), AI will reach its full potential only if, as part of its intelligence, it also has enough

teaming intelligence to work well with people. Although seemingly counterintuitive, the more intelligent the technological system, the greater the need for collaborative skills.

4. Allocation: reality shows that tasks are always interdependent, and humans and machines/AI always need to cooperate. When tasks are divided into parts, the interdependencies are frequently overlooked. The easiest subtask is then automated and the other subtasks are ignored. The moment the machine can no longer perform its subtask, as surprise is continuous, control is suddenly transferred to a human being who then experiences an ‘automation surprise’.
5. Workload and productivity: according to the Law of Stretched Systems (Woods & Hollnagel, 2006), automation is always exploited fully, requiring people to do more, do it faster, or in more complex ways, thereby increasing rather than decreasing workload. Also, new types of cognitive work are being created, often at the wrong moments (‘clumsy automation’), which leads to new types of errors.

This discussion on research on joint cognitive systems has prepared us for a discussion of how AI could enhance (or degrade) human expertise in various settings. In this next section, we illustrate the general principles we have described through a case study.

Case Study: Radiology

The modern work practice of radiology involves several healthcare professions working together as a team. A radiologist is a medical doctor who interprets medical images, communicates these findings to other physicians, and performs medical procedures using imaging. The radiographer produces medical images for the radiologist to interpret. The nurse is involved in patient care before and after imaging or procedures. It is clear that teamwork is vital, with a lot of interdependencies between various healthcare professions. Also, a variety of imaging techniques are used: radiographs (X-ray imaging), ultrasound, computed tomography, magnetic resonance imaging, and nuclear medicine. Each of these techniques requires specific expertise in terms of preconditions for use and sensitivity of data.

Radiological expertise not only involves a substantial perceptual component, but also involves the integration of several distinct bodies of knowledge with separate organizing principles, including physiology, anatomy, medical theories of disease, and the projective geometry of radiography (Lesgold et al., 1988). Lesgold and colleagues found that expert radiologists, when examining radiographs, would quickly (within two seconds) invoke a diagnosis schema that has prerequisites or tests that must be satisfied before it can control the diagnosis and viewing. The patient's anatomy is constructed as the schemata are applied. The expert works efficiently to reach the stage where an appropriate general schema is in control. When a schema does not fit the data, it is discarded quickly. On the other hand, schemata also drive perception by setting hypotheses on what to expect in an image. Each schema contains a set of processes that allows the viewer to reach a diagnosis and confirm it. The expert works both bottom-up, data-driven, as well as top-down, schema-driven, in a continuous cycle. This confirms the general picture of expertise we outlined, as it includes both recognitional decision making, based on a large and diverse memory for exemplars (Norman, Coblenz, Brooks, & Babcock, 1992), and being flexible and adaptive, but with more resource-intensive reasoning components, with the latter being employed in more difficult cases (Patel, Kaufman, & Kannampallil, 2020).

Over the last decade, modern AI technologies (particularly deep learning) have caused breakthrough successes in almost all areas of AI-assisted radiology. Examples include detecting and segmenting lung cancer tumors in radiographs, interpretation of MRI scans, and monitoring disease progress. For some of these tasks, AI achieved human level performance or better (Hosny, Parmar, Quackenbush, Schwartz, & Aerts, 2018). Despite the wide range of opportunities, these systems have not yet been implemented in clinical radiology practice. The earliest applications can be expected in areas where abundant high-quality-labeled data are available and concern tasks that currently overload human experts (such as in the detection of tumors in radiographs).

It is becoming clear that the introduction of this type of AI automation will not replace humans, but rather will lead to new workflows and create new roles for humans, requiring different human expertise. We can

expect the following types of task-changes in an AI-assisted radiology workflow:

AI-replacement tasks that are completely taken over by the AI. These are subtasks in the radiology workflow at which the AI consistently performs as well as or better than humans. Examples are the visual interpretation of radiology images by deep learning image classifiers. This will result in deskilling of existing radiology personnel and relieve training requirements of new radiologists from having to acquire this skill.

AI-augmentation tasks where the AI system augments humans. These are tasks for which the AI (e.g., due to brittleness) sometimes makes mistakes that can be repaired by humans. An example of this is planning a patient's treatment. Whereas AI can help in monitoring the effects of past treatments, it is highly unlikely that a treatment plan is finalized without any human oversight. For these tasks humans are needed to recognize and deal with abnormal and rare cases. This requires that humans maintain the expertise of this task and acquire additional expertise on how to use the AI support system. Also, the AI support system must be capable of explaining its advice to humans such that they can judge its trustworthiness.

AI-maintenance tasks are added to the radiology workflow that did not exist before. These have to do with maintaining the AI systems, and require a whole new set of skills. Examples of these are (re-)training the AI system, understanding the complications of introducing new hardware on the AI performance, training human personnel to use the AI system, and so on.

Despite the fact that humans will not be fully replaced, the efficiency (in terms of human labor) of radiology practice will undoubtedly increase due to the introduction of AI. However, these efficiency gains may well be compensated by a higher standard of healthcare, such as requiring more frequent health checkups.

Summarizing the trends in radiology, we can see that the hybrid AI principle, where different forms of AI work together with experts is the most appropriate vision. These different forms of collaboration will require different skills. Daugherty and Wilson (2018) refer to these skills as fusion skills, as they draw on the fusion of human and machine talents within a business process to create better outcomes than working

independently. For instance, *rehumanizing time* is a fusion skill that allows people to skillfully redirect their time toward more human activities. Particularly in medicine, physicians could greatly benefit from AI taking over visual interpretation of radiology images, as it would give them more time to see their patients or coordinate with other physicians. Other fusion skills involve knowing how best to ask questions of AI to get the insights you need, the ability to develop robust mental models of AI agents to improve process outcomes, and the ability to decide a course of action when a machine is uncertain about what to do.

Conclusion

Our review of the history of expertise studies with the history of AI has converged on a number of common themes. First, expertise is currently viewed as a skilled adaptation to complexity and novelty. This is not to diminish the importance of pattern-recognition capabilities amassed during extensive periods of deliberate practice. Rather, it is recognized that adaptation to complexity and novelty can only be skilled as a result of extensive practice. Second, although the current interest in AI largely focuses on machine learning capabilities there are a number of problems associated with that approach. First, machine learning approaches using deep neural networks cannot explain themselves to humans. This is crucial, particularly when experts need to work with these systems. Second, these approaches result in brittle systems that can easily be attacked or that do not work in unforeseen scenarios.

The history of joint cognitive systems has shown that viewing machines as prostheses results in breakdowns in performance, whereas viewing machines as tools or instruments aids in adapting to unanticipated variability. We have argued for a future of Hybrid AI in which expertise will be distributed across experts and AI in various ways. The example of radiology has shown that the introduction of AI capabilities may have various consequences, ranging from replacement, to augmentation, to maintenance of human expertise. It may well be the case that pattern recognition capabilities of AI systems will exceed human expertise (they already do so in restricted task domains). Yet, in order to be able to

effectively collaborate with human experts, AI will need collaborative skills, such as being able to explain itself to human experts. This is an area that is still being researched. Simultaneously, human flexibility and adaptation will increasingly be required to deal with unanticipated variability and surprise situations. Human expertise will be needed to close the gap between the demonstration and the real thing (Woods, 2016). This is in line with recent views on expertise that stress skilled adaptation to complexity and novelty.

Finally, the introduction of AI will also result in a whole series of new skills that human experts need to develop in order to deal with AI. We have discussed a few of these fusion skills (Daugherty & Wilson, 2018), but there are likely to be many more that we cannot foresee. AI systems will hardly ever stand alone in a work process and will therefore need intricate tuning to human demands at various points in time. Such systems will need to be trained, validated, understood, explained, assisted, and overruled if experts want to accept them and be able to effectively work with them.

This chapter has shown that it is a gross oversimplification to consider AI systems and human expertise as two mutually exclusive entities, with one taking over the other without changing anything in the work process. Rather, we need to view this from a joint cognitive systems perspective, at a systems level and as dynamically changing over time. Only then will we be able to see the intricacies of the mutual dependencies between humans and AI, and the constantly evolving distribution of skill sets that are required from an organizational perspective. There is a bright future for experts working jointly and collaboratively with AI systems in organizations.

References

- Anderson, J. R. (1981). *Cognitive skills and their acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779.
- Billings, C. (1991). *Human-centered aircraft automation: A concept and guidelines*. NASA Tech Memo 103885. Springfield, VA: National technical information service.

- Bobrow, D. G. (1984). Qualitative reasoning about physical systems: An introduction. *Artificial intelligence*, 24(1–3), 1–5.
- Bostrom, N. (2017). *Superintelligence*. Paris: Dunod.
- Bradshaw, J. M., Hoffman, R. R., Johnson, M., & Woods, D. D. (2013). The seven deadly myths of “autonomous systems”. *IEEE Intelligent Systems*, 28(3), 54–61.
- Brooks, R. A. (1990). Elephants don't play chess. *Robotics and autonomous systems*, 6(1–2), 3–15.
- Brown, N., & Sandholm, T. (2018). Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374), 418–424.
- Campbell, M., Hoane, A. J., Jr., & Hsu, F. H. (2002). Deep blue. *Artificial Intelligence*, 134(1–2), 57–83.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4, 55–81.
- Chen, Y., Elenee Argentinis, J., & Weber, G. (2016). IBM Watson: How cognitive computing can be applied to big data challenges in life sciences research. *Clinical Therapeutics*, 38(4), 688–701.
- Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121–125.
- Chi, M. T. H., Glaser, R., & Farr, M. J. (1988). *The nature of expertise*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Christoffersen, K., & Woods, D. D. (2002). How to make automated systems team players. In E. Salas (Ed.), *Advances in human performance and cognitive engineering research* (Vol. 2, pp. 1–12). Kidlington, UK: Elsevier Science.
- Daugherty, P. R., & Wilson, H. J. (2018). *Human + Machine: Reimagining work in the age of AI*. Cambridge, MA: Harvard Business Review Press.
- Davis, R. (1984). Amplifying expertise with expert systems. In P. H. Winston & K. A. Prendergast (Eds.), *The AI Business: commercial uses of artificial intelligence* (pp. 17–40). Cambridge, MA: MIT Press.
- de Groot, A. D. (1946). *Het denken van den schaker*. Amsterdam: Noord Hollandsche.
- de Groot, A. D. (1965). *Thought and choice in chess* (first Dutch edition in 1946). Mouton Publishers.
- Ericsson, K. A. (2006). An introduction to *The Cambridge Handbook of Expertise and Expert Performance*: Its development, organization, and content. In K. Anders Ericsson, N. Charness, P. J. Feltovich, & R. R. Hoffman (Eds.), *The Cambridge handbook of expertise and expert performance* (pp. 3–19). New York: Cambridge University Press.

- European Commission. (2019, April 8). *A definition of AI: Main capabilities and disciplines*. Independent high-level expert group on artificial intelligence. Brussels: European Commission. Retrieved from <https://ec.europa.eu/digital-single-market/en/news/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines>
- Feigenbaum, E. A. (1989). What hath Simon wrought? In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon* (pp. 165–182). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Feigenbaum, E. A., McCorduck, P., & Nii, H. P. (1988). *The rise of the expert company*. New York: Times Books.
- Feltovich, P. J., Prietula, M. J., & Anders Ericsson, K. (2006). Studies of expertise from psychological perspectives. In K. Anders Ericsson, N. Charness, P. J. Feltovich, & R. R. Hoffman (Eds.), *The Cambridge handbook of expertise and expert performance* (pp. 41–67). Cambridge, UK: Cambridge University Press.
- Ferrucci, D. A. (2012). Introduction to “This is Watson”. *IBM Journal of Research and Development*, 56. (3.4), 1, 1–1:15.
- Germain, M.-L. (2006, February). *Stages of psychometric measure development: The Example of the Generalized Expertise Measure (GEM)*. Paper presented at the Academy of Human Resource Development International Conference, Columbus, OH. Retrieved from <http://www.eric.ed.gov/PDFS/ED492775.pdf>
- Germain, M.-L., & Tejada, M. J. (2012). A preliminary exploration on the measurement of expertise: An initial development of a psychometric scale. *Human Resource Development Quarterly*, 23, 203–232.
- Glaser, R., & Chi, M. T. H. (1988). Overview. In M. T. H. Chi, R. Glaser, & M. Farr (Eds.), *The nature of expertise* (pp. xv–xxviii). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gobet, F. (2020). The classic expertise approach and its evolution. In P. Ward, J. M. C. Schraagen, J. Gore, & E. M. Roth (Eds.), *The Oxford handbook of expertise* (pp. 35–55). Oxford, UK: Oxford University Press.
- Goldstein, I., & Papert, S. (1977). Artificial intelligence, language, and the study of knowledge. *Cognitive Science*, 1(1), 84–123.
- Gottfredson, L. S. (1997). Mainstream science on intelligence: An editorial with 52 signatories, history, and bibliography. *Intelligence*, 24, 13–23.
- Grenier, R. S., & Germain, M.-L. (2014). Expertise through the HRD lens: Research trends and implications. In N. F. Chalofsky, T. S. Rocco, & M. L. Morris (Eds.), *Handbook of Human Resource Development* (pp. 183–200). Hoboken, NJ: John Wiley & Sons.

- Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2), 199–220.
- Gunning, D., & Aha, D. W. (2019). DARPA's explainable artificial intelligence program. *AI Magazine*, 40(2), 44–58.
- Hambrick, D. Z., Burgoyne, A. P., & Oswald, F. L. (2020). Domain-general models of expertise: The role of cognitive ability. In P. Ward, J. M. C. Schraagen, J. Gore, & E. M. Roth (Eds.), *The Oxford handbook of expertise* (pp. 56–84). Oxford, UK: Oxford University Press.
- Hatano, G., & Inagaki, K. (1986). Two courses of expertise. In H. Stevenson, H. Azuma, & K. Hakuta (Eds.), *Child development and education in Japan* (pp. 262–272). New York: W. H. Freeman.
- Holyoak, K. (1991). Symbolic connectionism: toward third-generation theories of expertise. In K. Anders Ericsson & J. Smith (Eds.), *Toward a general theory of expertise: Prospects and limits* (pp. 301–335). Cambridge, UK: Cambridge University Press.
- Hosny, A., Parmar, C., Quackenbush, J., Schwartz, L. H., & Aerts, H. J. (2018). Artificial intelligence in radiology. *Nature Reviews Cancer*, 18(8), 500–510.
- Hunt, E. (2006). Expertise, talent, and social encouragement. In K. Anders Ericsson, N. Charness, P. J. Feltovich, & R. R. Hoffman (Eds.), *The Cambridge handbook of expertise and expert performance* (pp. 31–38). Cambridge, UK: Cambridge University Press.
- Johnson, J. G., & Raab, M. (2003). Take the first: Option-generation and resulting choices. *Organizational Behavior and Human Decision Processes*, 91, 215–229.
- Johnson, M., & Vera, A. (2019). No AI is an island: The case for teaming intelligence. *AI Magazine*, 40(1), 16–28.
- Klein, G., Wolf, S., Militello, L., & Zsombok, C. (1995). Characteristics of skilled option generation in chess. *Organizational Behavior and Human Decision Processes*, 62(1), 63–69.
- Klein, G., Woods, D. D., Bradshaw, J., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a “team player” in joint human-agent activity. *IEEE Intelligent Systems*, 91–95.
- Leith, P. (2016). The rise and fall of the legal expert system. *International Review of Law, Computers & Technology*, 30(3), 94–106.
- Lesgold, A., Rubinson, H., Feltovich, P., Glaser, R., Klopfer, D., & Wang, Y. (1988). Expertise in a complex skill: Diagnosing x-ray pictures. In M. T. H. Chi, R. Glaser, & M. Farr (Eds.), *The nature of expertise* (pp. 311–342). Hillsdale, NJ: Lawrence Erlbaum Associates.

- Mamede, S., Schmidt, H. G., Rikers, R. M. J. P., Custer, E. J. F. M., Splinter, T. A. W., & Van Saase, J. L. C. M. (2010). Conscious thought beats deliberation without attention in diagnostic decision-making: At least when you are an expert. *Psychological Research*, *74*, 586–592.
- Marcus, G., & Davis, E. (2019). *Rebooting AI: Building artificial intelligence we can trust*. New York: Random House.
- Minsky, M., & Papert, S. A. (1972). *Artificial intelligence progress report*.
- Moosavi-Dezfooli, S. M., Fawzi, A., & Frossard, P. (2016). Deepfool: A simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2574–2582).
- Moxley, J. H., Anders Ericsson, K., Charness, N., & Krampe, R. T. (2012). The role of intuition and deliberative thinking in experts' superior tactical decision-making. *Cognition*, *124*, 72–78.
- Nilsson, N. J. (2009). *The quest for artificial intelligence*. Cambridge, UK: Cambridge University Press.
- Norman, G. R., Coblenz, C. L., Brooks, L. R., & Babcock, C. J. (1992). Expertise in visual diagnosis: A review of the literature. *Academic Medicine*, *67*(10), S78–S83.
- Patel, V. L., Kaufman, D. R., & Kannampallil, T. G. (2020). Diagnostic reasoning and expertise in healthcare. In P. Ward, J. M. Schraagen, J. Gore, & E. Roth (Eds.), *The Oxford handbook of expertise* (pp. 618–641). Oxford, UK: Oxford University Press.
- Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. New York: Basic Books.
- Peeters, M. M. M., van Diggelen, J., van den Bosch, K., Bronkhorst, A., Neerincx, M. A., Schraagen, J. M., & Raaijmakers, S. (2020). Hybrid collective intelligence in a human-AI society. *AI & Society*. <https://doi.org/10.1007/s00146-020-01005-y>
- Quillian, R. (1963). *A notation for representing conceptual information: An application to semantics and mechanical English para-phrasing*. SP-1395, System Development Corporation.
- Roth, E. M., Bennett, K. B., & Woods, D. D. (1987). Human interaction with an “intelligent” machine. *International Journal of Man-Machine Studies*, *27*, 479–525.
- Sarter, N., & Woods, D. D. (1995). “How in the world did we get into that mode?” Mode error and awareness in supervisory control. *Human Factors*, *37*, 5–19.

- Schraagen, J. M. C. (2018). Naturalistic decision making. In L. J. Ball & V. A. Thompson (Eds.), *The Routledge international handbook of thinking and reasoning* (pp. 487–501). London: Routledge.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., & Hassabis, D. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354–359.
- Simon, H. A., & Gilmartin, K. (1973). A simulation of memory for chess positions. *Cognitive Psychology*, 5, 29–46.
- van Harmelen, F., & Teije, A. T. (2019). A boxology of design patterns for hybrid learning and reasoning systems. *arXiv preprint arXiv:1905.12389*.
- Vergne, M. (2017). Artificial Intelligence and expertise: The two faces of the same artificial performance coin. In *Proceedings of the AAAI-17 Workshop on Human-Machine Collaborative Learning* (pp. 696–702).
- Ward, P., Gore, J., Hutton, R., Conway, G., & Hoffman, R. (2018). Adaptive skill as the *conditio sine qua non* of expertise. *Journal of Applied Research in Memory and Cognition*, 7(1), 35–50.
- Ward, P., Schraagen, J. M. C., Gore, J., Roth, E. M., Hoffman, R. R., & Klein, G. (2020). Reflections on the study of expertise and its implications for tomorrow's world. In P. Ward, J. M. C. Schraagen, J. Gore, & E. M. Roth (Eds.), *The Oxford handbook of expertise* (pp. 1193–1213). Oxford University Press.
- Wiener, E. L. (1989). *Human factors of advanced technology ("glass cockpit") transport aircraft* (NASA Contractor Report No. 177528). NASA Ames Research Center.
- Woods, D. D. (2016). The risks of autonomy: Doyle's catch. *Journal of Cognitive Engineering and Decision Making*, 10(2), 131–133.
- Woods, D. D., Dekker, S., Cook, R., Johannesen, L., & Sarter, N. (2010). *Behind human error* (2nd ed.). Farnham, UK: Ashgate Publishing Limited.
- Woods, D. D., & Hollnagel, E. (2006). *Joint cognitive systems: Patterns in cognitive systems engineering*. Boca Raton, FL: CRC Press.
- Woodworth, R. S., & Schlosberg, H. (1954). *Experimental psychology* (Rev. ed.). New York: Henry Holt and Company.