

Editor's Review and Introduction:
Cognition-Inspired Artificial Intelligence

Daniel N. Cassenti*

Computational & Information Sciences Directorate
DEVCOM Army Research Laboratory

Vladislav D. Veksler

School of Business and Computer Science
Caldwell University

Frank E. Ritter

College of Information Sciences and Technology
Pennsylvania State University

Keywords: Cognitive Science; Artificial Intelligence; Fitts' List; Human-AI Collaboration; Cognitive Systems

* Corresponding Author Information

FCDD-RLC-CA
U.S. Army Research Laboratory
2800 Powder Mill Rd
Adelphi, MD 20783
1-301-787-3695
daniel.n.cassenti.civ@army.mil

Abstract

Cognitive science has much to contribute to the general scientific body of knowledge, but it is also a field rife with possibilities for providing background research that can be leveraged by artificial intelligence (AI) developers. In this introduction, we briefly explore the history of AI. We particularly focus on the relationship between AI and cognitive science and introduce this special issue that promotes the method of inspiring AI development with the results of cognitive science research.

1. Introduction

The field of cognitive science today is often relegated to academic concerns. There is a wide range of scientific questions and concerns that can be addressed by cognitive scientists. From sensation and perception to problem solving, animal and human cognition has many unexplored and unanswered puzzles and mysteries.

As cognitive scientists figure out how human intelligence is generated by mental processes, artificial intelligence (AI) developers can use this knowledge to develop AI rather than trying to create intelligence from completely introspective claims and assumptions about how it should be composed. Knowledge of human cognition should influence AI development in principle; however, the literature indicates that this influence has occurred too infrequently. In this special issue, we offer a set of research articles that argue and provide examples as to how cognition can aid and influence AI system design.

This special issue has many similarities to Leito and Radicioni’s special issue (2016a), though our special issue covers a larger variety of cognitive and AI research, in more depth, and with greater adherence to how cognitive research inspires AI (e.g., two of their five articles do not use the term “artificial intelligence”). Additionally, we propose a radical shift in the way both the cognitive science and AI communities view the relationship between the two fields, which is not attempted in Leito and Radicioni’s (2016b) introductory article.

Before we delve into the articles, it is important to understand how AI was first conceptualized and what it has become. The link between human and artificial intelligence has not been lost on some of the greatest minds in cognitive science, including Alan Turing, Herb Simon, and Allen Newell. We will review the history of the relationship between cognitive science and artificial intelligence and show how early cognitive scientists understood how AI could help cognitive science, but not necessarily the reverse. Only recently have more researchers sought ways in which cognitive science could aid the development of AI.

2. What is Artificial Intelligence?

Artificial intelligence is any programming in a computerized device that emulates human higher-level cognitive abilities. Higher-level cognitive abilities include such things as pattern recognition, learning, and problem solving; whereas lower-level cognitive abilities such as

memory storage or executing prescribed if-then algorithms (e.g., as insects do, see Theraulaz & Bonabeau, 1995) do not make a program intelligent.

Buchanan (2006) has a very good article reviewing the history of AI with much of the following historical information also reported in his review. As often occurs with various innovative concepts, the original idea of AI came from philosophical thought with Descartes describing a mechanical man and Pascal designing machines that could perform arithmetic calculations. The Mechanical Turk inventors purported to have the first thinking automaton when they introduced their device in 1770 and claimed that it could play chess without human intervention (though the device contained a human inside who controlled the puppet figure).

Alan Turing established the basis for AI with his conceptualization of a machine with limitless symbolic memory and an internal scanner that could retrieve information in 1935. This turned out to be the basis for computerized devices up to the current time. In 1951, Christopher Strachey created the first AI program that could play checkers, followed closely by Dietrich Prinz’s AI program that could play chess. A year later, Anthony Oettinger created the first AI program that could learn by creating his Shopper program, which visited virtual stores, remembered items the store contained, and went directly to them when those items were required instead of randomly searching stores on subsequent trials. These may seem like simple programs now, but for the time, they were revolutionary steps forward in the evolution of AI.

Also in 1951, Paul Fitts (1951) led a panel to describe the skills machines perform better than humans and vice-versa. The most famous outcome of the panel was a list of statements pertaining to the comparative strengths of machines and humans, which became known as “Fitts’ List.” Fitts’ List is summarized in Table 1. Although there have been plenty of disagreements with Fitts’ List, it has been fairly robust to critique in the literature (see Cummings, 2014) and in a recent survey of engineering students (de Winter & Hancock, 2015).

Table 1. Summarized Fitts’ List.

Humans Better Than Machines	Machines Better Than Humans
Detecting inconsistencies	Operating in extreme environments
Diversity of stimuli sensations	Sensing particular stimuli
Perceiving patterns	Response speed
Attention to relevancy	Processing speed
Creative thinking	Precision in repetition
Strategic task allocation when overloaded	Multi-tasking
Flexibility	Smooth force exertion
Learning from experience	Accurate performance
Handling of low-probability events	Impervious to distraction
Induction	Deduction

Today, AI has become much more complex, though it retains the basic principles of Fitts’ List of doing better with the items in the right column than the left column of Table 1. AI can search

databases for information with incredible speed and accuracy relative to humans. It can perform multiple tasks at once without getting distracted or overwhelmed. It has the ability to use particular stimuli to perform the same operations without regard for its personal circumstances. AI developers have creatively used these skills to ramp up the evolution of AI, yet they rarely take AI out of the realm of the right-hand side of Fitts' List.

However, there is no conclusive reason that would prohibit AI from making strides in reproducing the left-hand side of Fitts' List. One could cite neuronal plasticity and volume to argue that humans have an advantage over AI, but neuromorphic devices are now a reality (see Tang, Yuan, Shen et al., 2018), allowing AI to use neuronal cells and thus retain the plasticity of neurons. As for volume, Moore's (1965) Law claims that the number of transistors on a chip doubles every two years. The Law has been fairly robust over time but is going to end sometime this decade (see Rotman, 2020). If Moore's Law continues until 2026, Peper (2017) claims that the number of transistors on a circuit board would surpass the number of neurons in the human brain (i.e., ~100 billion). Though this is not a foregone conclusion, clearly volume is not as much of a limiting factor for achieving the power of the human brain as it might seem.

The field of study that aims to understand human intelligence is cognitive science. Our position is that the best way to develop AI that achieves the skills that humans possess involves research in cognitive science. If a developer can understand how cognitive scientists understand and model human intelligence, then they can emulate the model in their AI programming and attempt to achieve the same.

3. Early AI researchers recognized the link between human and artificial intelligence

The Turing Test, initially introduced as the Imitation Game (Turing, 1950) is the most well-known thought experiment involving AI. According to Alan Turing, we would know a computer could truly think if it could be mistaken for a human in a conversation imitating another human. Underlying this formulation is the idea that thinking is inherently human and truly thinking machines produce behavior that could be mistaken for that of humans. We can consider this idea as a precursor to the position of this special issue that human intelligence is a model for AI and indeed represents a sort of benchmark to indicate when AI achieves new milestones.

Since Turing published his paper, the Turing Test has been taught in psychology, computer science, and philosophy classes. Given that, there is an underlying recognition that success with AI depends on replication of human thought patterns. It would stand to reason that studying cognitive science would be a positive training opportunity to help drive creation of new AI capabilities. This is not yet a standard practice in colleges or universities to the present day.

Nobel-Prize winner Herbert Simon recognized that psychology drove early AI conceptualization (see Simon, 1977). According to Simon, early AI researchers borrowed terms from psychology to drive the creation of the field. There appeared to be an initial impetus to accept psychology as the basis for AI development, yet the field has not generally followed along that course as development continued.

A possible reason that AI researchers backed away from the idea that psychology should form the basis of AI is that the cognitive revolution of the mid-to-late 1950's did not occur until after the initial conceptualizations of AI. Borrowing terms from psychology in the 1950s would not have been that helpful for AI development given the predominant school of thought for scientific study of psychology at that time was behaviorism (see Gardner, 1987 for a review of the history of cognitive science, including the role of behaviorism). Behaviorism framed psychological study as an attempt to catalogue and understand stimulus-response pairings and thereby disregard the processing that led to those responses as belonging to a black box (i.e., cognitive processing is unknowable because it cannot be directly observed, see Ryan, 2019). Psychology was not restricted to behaviorism, however, movements such as Gibson's ecological psychology (Lobo, Heras-Escribano, & Travieso, 2018) similarly did not posit much regarding internal processing of sensation-action pairings, nor did clinical psychology offer much when the field focused on psychological dysfunction that could not help AI researchers understand how to build functional AI.

In 1956, George Miller's (1956) paper, "The Magic Number Plus or Minus Two," introduced the concept of short-term-memory capacity limitations. This was arguably the first cognitive science paper and created ripples in the psychological-science community as it broke from the pervading behaviorist stance. By then, the advancements in AI listed in the previous section had already been achieved. In addition, the term "Artificial Intelligence" had already been coined by John McCarthy in 1955 (see McCarthy, Minsky, Rochester, & Shannon, 2006). It is understandable that cognitive science did not help AI developers given that AI as a field came first. Indeed, early thinking about the relationship between AI and cognitive science was not our position that cognitive science can act as the basis for developing AI, but that AI can advance the field of cognitive science. We will discuss this concept in the next section.

4. AI as a means of advancing the field of cognitive science

When examined further, the predominance of the idea that AI would help cognitive science more than the reverse fits within the context of where both fields were when they began only a few years apart. The Miller (1956) article was not the death knell for behaviorism at which point all psychologists abandoned it in favor of cognitive psychology. Even the dismantling of behaviorism by Noam Chomsky (1959) did not eliminate behaviorism from its predominant state. Cognitive psychology needed to first break down the barriers put up by behaviorism before it could be accepted as a dominant approach. Cognitive psychology faced resistance and push back throughout its early development for its attempt to examine the black box, which arguably stalled development from the late 1950's through the 1970's.

As a field, AI began brand new. It did not meet resistance from what was there before it, but instead was permitted to flourish and gave AI developers the opportunity to flexibly explore the avenues of creative problem solving. Newell (1970) sought to capitalize on this unfettered exploration by writing a paper arguing for using AI to help advance cognitive science. We can see his frustration with the state of cognitive science even in the late 1970's when he suggests at the end of the paper that behaviorism had stigmatized the study of thinking and problem solving so much that he did not have much optimism for the future of scientific investigation in those

areas. Clearly, in Newell's assessment, AI had over two decades of fluid development and cognitive science had made less progress.

Newell's (1970) argument centered on the notion that in the creation of skills, AI had shown the ability to construct cognitive skills that allowed it to operationalize those behaviors. AI developers created these skills by programming them and thus leave behind code that can be examined. This essentially broke down the barrier asserted by behaviorists of the inability to observe mental concepts when AI code is directly observable.

Newell's perspective ushered in the popular analogy of the mind as a computer. Newell's (1970) understanding of mental concepts was operationalized by examining how AI could use symbolic processing to achieve thinking and problem solving and analogizing that to human mental processing.

Newell and Simon (1972) had pioneered the bridge between AI and cognition with a style of programming known as a production system, the first type of cognitive modeling system. A production system operated through symbolic transformations using if-then statements where a certain set of conditions would trigger a certain set of actions. They called this system the General Problem Solver and it was the precursor to the most widely used production systems, ACT-R (Anderson, 2007; Anderson & Lebiere, 1998; Ritter, Tehranchi, & Oury, 2019) and Soar (Newell, 1990; Laird, 2012).

Furthermore, Simon (1981) took this a step further and wrote about how computer simulations can help us understand mental processing. This perspective opened up a new field of cognitive modeling that uses symbolic and sub-symbolic processing to reproduce the results of empirical studies. If the modeler can create a model that replicates the experimental results, then the model may be seen as analogous to how human participants generated the same results.

Cognitive tutoring (see Anderson & Gluck, 2001) is an AI application that helps humans and arose out of cognitive modeling. With the understanding of how cognition works, AI can monitor learners and adjust the speed and complexity of their learning process by judging performance and adjusting to teach those component cognitive skills that the learner needs to develop before mastering the subject area.

As another example, adaptive automation (e.g., Kaber, Wright, Prinzel, & Clamann, 2005) is an application of artificial intelligence that can help humans by understanding how cognition works. Adaptive automation is AI that monitors a computer user's activity for signs of difficulty performing a task. If the adaptive automation does not detect any difficulties, then it does nothing so as not to interrupt good performance. If there are signs of difficulty, the adaptive automation will attempt to aid the user with an intervention to boost performance.

Cassenti, Gamble, and Bakdash (2016) described four ways that adaptive automation can detect difficulty with a task. The AI could rely on the user to determine when they need help by triggering an aid when the user takes an action to activate the aid. It could also gauge some form of ongoing performance measure and intervene when performance dips below a certain threshold. A third way would be to measure physiological signs of problem, such as pupil

diameter indicating high workload (Chen & Epps, 2014). A final way is to use cognitive modeling to determine a priori when the most difficult times will be when completing a task and intervene at those times. Any combination of these methods could be implemented as well.

Cognitive modelers today continue to practice both cognitive and computer science, and as such represent the best chance to understand the benefit AI can get from cognitive science. The editors and corresponding authors of this special issue are all cognitive modelers, who have skills in analyzing cognition; particularly how cognitive mechanisms process information to generate behavior and how they link together to go from an initial state to a goal state. Thus, for now, cognitive modelers are the most likely of all professionals to be able to apply cognition to create new AI skills.

Over time, there has been an ebb and flow of ideas and inspiration between these two fields. In the next section, we will review how researchers are beginning to see how cognitive science can once again help AI.

5. Only recently have researchers discussed how cognitive science can help AI

In the 1980s, cognitive science finally displaced behaviorism as the primary scientific discipline studying human behavior. This came about largely by efforts from cognitive modelers such as Herbert Simon, Alan Newell, and John Anderson who argued for the computer metaphor of the mind, which hinged on the notion of cognition as information processing. Seen from this angle, cognitive science owes its accomplishment of overcoming behaviorism to computer science and more specifically to AI for the reasons described in the previous section. Without advancements in AI that showed creative ways to simulate intelligent behavior, the argument that we can examine computer code and therefore overcome the black box would not have been accepted because computer code and cognition would have been considered distinct phenomena.

The best example of how cognitive science has helped AI development is perhaps Parallel Distributed Processing (PDP; Rumelhart, Hinton, & McClelland, 1986; McClelland Rumelhart, & Hinton, 1986), colloquially known as connectionism. Connectionist modeling showed that computers could learn to classify stimuli with the proper configurations of nodes and connections between them. This type of modeling was inspired by the biology of the brain with nodes representing neurons and connections representing dendrites. Unfortunately, as McClelland and Rumelhart's ideas became prominent in cognitive science, AI was beginning what is known as the AI Winter (see McClay, 1995), which was a lull in interest in AI that hit a peak in the early 1990s. It took until the early 2000s for machine learning to become prominent in the field of AI and by then interest among cognitive scientists in PDP had faded (see Bengio, Goodfellow, & Courville, 2016 for more details on the history of machine learning). However, even if the evolution of machine learning is not always acknowledged, there is no doubt that connectionism helped create machine learning, the most prominent field of AI at the time of publication of this special issue.

Bayesian statistics is another topic of study within the discipline of cognitive science that contributed to AI development. Kahneman and Tversky (1973) first understood that human

decision making suffered from base-rate neglect with considerations of base rate as the key statistical foundation for decisions. Gigerenzer et al. (2007) followed up on this finding to argue that neglecting base rates when interpreting medical-test evidence is the primary driver in misdiagnosis of medical conditions. The article had widespread impact particularly in cognitive psychology, but also in AI (e.g., Allen, Singh, Greiner, & Hooper, 2008) in which subjective logic depends on Bayesian statistics to aid AI to make decisions under uncertainty (Jøsang, 2016).

Six years later, Langley (2012) advocated for his cognitive systems perspective, perceiving a departure of AI from its original focus on developing human higher-level cognitive functions. In his view, AI developers had switched attention to the vast storage capacity and fast processing speeds of computers and lost sight of the recreation of intelligent cognitive processing. In the first issue of the journal, *Advances in Cognitive Systems*, Langley built his argument for the approach, leading a group of essays addressing some aspects of his cognitive systems perspective.

These essays included one where Fahlman (2012) argued that AI researchers lost their way by ignoring human higher-level cognition to focus on creating brittle AI programs to solve narrow problems. Cassimatis (2012) asked the reader to reject the old AI scientific approach and develop a new one. Forbus (2012) described how the field of AI has gone out of balance and plots out a course for returning to a focus on building minds. Bello (2012) advocated for mind-reading skills in AI. Sammut (2012) discussed how human cognition can help design smarter robots. Lastly, Jones, Wray, and van Lent (2012) argued for evaluative criteria and requirements for AI based on cognitive science.

Although we applaud Langley's efforts and appreciate the contribution his introduction to the first issue of *Advances in Cognitive Systems* made to tying cognitive science back in with AI, his vision is somewhat different from ours. Langley's vision is to return to the original intent of AI to ask the same questions as cognitive science does and specifically to build systems out of cognition. Our vision is to cover the spread of depth in cognitive research from simple phenomena to whole cognitive architectures to achieve new advances in AI. The AI researcher need not build up whole cognitive systems so long as new AI skills arise from something cognitive. The articles in this special issue cover this range of depth in cognitive science.

As an example, consider Kelley and Cassenti (2011) who show how AI in robots may be advanced through principles of cognitive developmental psychology. The authors review three areas: knowledge representation during stages of development; Levels of Consciousness Theory (Zelazo, 2000, 2004); and object permanence. They speculate about how each can help lead to advances in the way robots navigate, and how robotics can help further understanding of developmental psychology.

Our larger agenda in creating this special issue is to change the way in which we understand AI and to make the argument that cognitive science is integral to AI. We posit that AI is best considered an engineering discipline where a developer is attempting to create something new rather than a science that discovers what already exists. Behind every engineering discipline, there are one or more scientific disciplines that provide the basis for what the engineer is attempting to create. In the same way that mechanical and electrical engineers rely on their knowledge of physics to build new inventions, and in the same way that chemical engineers lean on knowledge of chemistry, we believe that artificial intelligence engineers should be knowledgeable in the foundations of cognitive science. This is already the case with human factors engineers, who often learn from coursework in cognitive science but it is still all too rare for artificial intelligence engineers to incorporate this foundational work into their curriculum.

It is reasonable to assume in light of the history of cognitive science that AI as a discipline did not develop a dependency on cognitive science, given that the latter was slower in its maturation. However, the study of mental processing, the activity that generates human intelligence, is now alive and well within cognitive science. It is again an opportune time for the developers of AI to take inspiration and direction from foundational theory and models in cognitive science.

An AI engineer without a background in cognitive science, which is most often the case, can rely on little to no knowledge about how intelligence works in the natural world. In social psychology, ascribing the cause of your own behavior to the environment and the cause of others' behaviors to their choices, is the fundamental attribution error. In development, ascribing how you think you will behave to others is the fundamental attribution error of design (Baxter, Churchill, & Ritter, 2014). Relying strictly on one's intuition, with little knowledge of established scientific theories and models has the potential to cause such attribution errors; at best this is inefficient, and at worst this may mislead AI engineering efforts altogether.

A culture shift needs to take place before we can create a situation where AI developers see themselves as engineers with cognitive science as the foundation of their knowledge. In our estimation, creating that shift requires at least three steps. First, a strong argument needs to be made in the literature where the case is made and backed up by multiple examples of cognition-inspired AI. This special issue is the attempt to make that argument. Next, we need to make inroads among the AI community to demonstrate the utility of cognitive science to that field. Lastly, we need to bring this argument to the academic community and convince those with influence over curricula for AI programs that students need to also study courses in cognitive psychology and cognitive science.

On this last point, changing the way that AI developers are educated constitutes the most important step in the three-step process, an approach also supported by Langley (2012). Cognitive science has been largely relegated to academic pursuits, rather than practical ones, although parts of human factors and HCI are exceptions to the rule. Yet, even if one were to examine the agendas of human factors conferences, such as the Human Factors and Ergonomics Society (HFES) Conference, much of the work done in this area presents the scientific, experimental, or modeling aspects of cognitive science rather than new creations and engineering achievements. Instead, the papers mostly speak to possible areas of application. This is not to

discount the contribution of human factors, but to accept that there are relatively few human factors programs at colleges or universities, so many of those presenting and publishing at conferences like HFES are still initially trained as cognitive psychologists or cognitive scientists and the papers reflect that expertise and those interests. The root cause of the state of human factors as a field is largely a reflection of the training received; and if not treated as an engineering field in its own right, human factors may end up having less impact than it could.

AI is not in the same place as human factors. It is a well-established field with many active AI programs in academic institutions around the globe. However, there is a deficit within these programs for supporting cognitive science in the curricula. In a list of the 25 best AI programs, Akins (2020) reports the classes and departments that support the programs. In only five of the 25 programs does cognitive science play a role. For something as fundamental as a core science-engineering relationship, a 20% rate of cognitive science education in AI is far from ideal. We aim to show through this introductory article and the contributed articles in this special issue that cognitive science is fundamental to the creation of artificial intelligence.

6. Contributions to this topic

The following are summaries of the topic contributions. The authors were selected and contacted on the basis of their status as cognitive modelers, a sub-field of cognitive science that represents the greatest overlap with computer science fields and thus most likely to address how cognition can develop AI. Subsequent discussions supported a wide range of cognitive phenomena with many insights for developing related AI systems to most efficiently explore the wide scope of how cognitive science can assist AI.

In “Theory of mind from observation in cognitive models and humans,” Ngoc Nguyen and Cleotilde Gonzalez of Carnegie Mellon University demonstrate how an existing cognitive learning theory may be used to develop AI agents as observers that infer and predict potential behavior of acting agents in a navigation task. The observer is able to pass a standard “Theory of Mind” test used in humans. Fully realized, this ability would increase AI understanding of human collaborators and improve human-machine teaming.

Sébastien Hélie of Purdue University writes in “When is psychology research useful in artificial intelligence? A case for reducing computational complexity in problem-solving” that human-brain processing is much slower than digital processing in computers yet the cognitive capacity to solve problems is often superior to AI. He discusses how a solution to this problem is to develop AI with comparable cognitive problem-solving efficiency and maps out how to achieve this.

In “Symbolic deep networks: A psychologically-inspired light-weight and efficient approach to deep learning” Vladislav Veksler, Blaine Hoffman, and Norbou Buchler of Caldwell University and the U.S. Army DEVCOM Data & Analysis Center review the efficiency and robustness of human learning processes and a variety of mechanisms in the psychological literature shown to

enable these features. The authors apply this knowledge to create a new type of deep learning that can advance the state of the art for AI development.

Sebastian Blum of Technical University Berlin and his colleagues discuss human-AI teaming in their paper, "Cognitive modeling of anticipation: Unsupervised learning and symbolic modeling of pilots' mental representations." According to the authors, human team members have some degree of ability to anticipate what joint actions their team could take to reach a goal state. They review the cognitive literature on anticipatory thinking and show how AI can be developed to create a similar skill, which could improve human-AI collaborative performance.

In "Cognition-enhanced machine learning models for better prediction with limited data," Florian Sense of the University of Groningen and his colleagues discuss how predictive analytics has been studied in the cognitive-science and AI-development communities but represent distinct approaches. The authors explore how to integrate cognitive-science into AI to form a new approach for predictive-analytic development in AI.

Konstantinos Mitsopoulos of Carnegie Mellon University and his colleagues discuss deep learning in "Towards a psychology of deep reinforcement learning agents using a cognitive architecture." They claim that deep learning is an AI methodology that is relatively detached from the study of human cognition. The authors apply a cognitive-architectures approach of how human reinforcement learning works to better explain deep learning in computers and how to use this information to improve AI.

The focus of "Knowledge gaps: A challenge for agent-based automatic task completion" by Goonmeet Bajaj of Ohio State University and her colleagues is on human-AI communication. They claim that when task knowledge is lacking, there are bound to be difficulties performing the task; an issue shared by both humans and AI. The authors outline what we know about the results of gaps in human task knowledge to address how to improve AI performance when task knowledge is lacking.

Lastly, Shashank Uttrani of the Indian Institute of Technology and colleagues discuss how cognitive principles of experiential memory shape human decision making to keep cognitive biases from influencing problem-solving decisions in "Life and death decisions and COVID-19: Investigating and modeling the effect framing, experience, and Context on preference reversals in the Asian disease problem." They show how AI can inform cognitive science through using a modeling framework to help create models quickly and using a genetic algorithm can help understand model-to-data comparisons. They also show that the path goes in the opposite direction in how these decision-making principles can shape AI problem-solving capabilities to support decision making.

This paper collection represents a continued impetus to enhance and extend the field of cognition-inspired AI. Continuing to inform the art of AI based on human cognition requires a set of research studies that demonstrate the ability to use the approach to make better AI that

Cassenti, D. N., Veksler, V. D., Ritter, F. E. (2022). Editor's Review and Introduction: Cognition inspired artificial intelligence. *Topics in Cognitive Science*. 14. 652-664.

helps to move AI across the division of Fitts' List. The collection of articles in this special issue extends progress in this direction.

7. Papers in this topic

Nguyen, T. N. & Gonzalez, C. Theory of mind from observation in cognitive models and humans.

Hélie, S. & Pizlo, Z. When is psychology research useful in artificial intelligence? A case for reducing computational complexity in problem-solving.

Veksler, V. D., Hoffman, B. E., & Buchler, N. Symbolic deep networks: A psychologically-inspired light-weight and efficient approach to deep learning.

Blum, S., Klaproth, O. W., Russwinkel, N. Cognitive modeling of anticipation: Unsupervised learning and symbolic modeling of pilots' mental representations.

Sense, F., Wood, R., Collins, M. G., Fiechter, J., Wood, A., Krusmark, M., Jastrzembski, T., Myers, C. W. Cognition-enhanced machine learning models for better prediction with limited data.

Mitsopoulos, K., Somers, S., Schooler, J., Lebiere, C., Pirolli, P., Thomson, R. Towards a psychology of deep reinforcement learning agents using a cognitive architecture.

Bajaj, G., Current, S., Schmidt, D., Bandyopadhyay, B., Myers, C. W., Parthasarathy, S. Knowledge gaps: A challenge for agent-based automatic task completion.

Uttrani, S., Sharma, N. & Dutt, V. Life and death decisions and COVID-19: Investigating and modeling the effect framing, experience, and Context on preference reversals in the Asian disease problem

8. References

Allen, T. V., Singh, A., Greiner, R., & Hooper, P., (2008). Quantifying the uncertainty of a belief net response: Bayesian error-bars for belief net inference, *Artificial Intelligence* 172, 483–513.

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: Oxford University Press. Anderson, J. R., & Gluck, K. (2001). What role do cognitive architectures play in intelligent tutoring systems. In S. M. Carver & D. Klahr (Eds.) *Cognition & instruction: Twenty-five years of progress*, (pp. 227-262). New York: Psychology Press.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.

- Cassenti, D. N., Veksler, V. D., Ritter, F. E. (2022). Editor's Review and Introduction: Cognition inspired artificial intelligence. *Topics in Cognitive Science*, 14, 652-664.
- Baxter, G., Churchill, E. F., & Ritter, F. E. (2014). Addressing the fundamental error of design using the ABCS. *AIS SIGHCI Newsletter*, 13(1), 9-10.
- Bello, P. (2012). Cognitive foundations for a computational theory of mindreading. *Advances in Cognitive Systems*, 1, 59-72.
- Bengio, Y., Goodfellow, I., & Courville A. (2017). *Deep learning*. Cambridge, MA: MIT Press.
- Buchanan, B. G. (2005). A (very) brief history of artificial intelligence. *AI Magazine*, 26, 53-53.
- Cassenti, D. N., Gamble, K. R., & Bakdash, J. Z. (2016). Multi-level cognitive cybernetics in human factors. In K. Hale and K. Stanney (Eds.) *Advances in neuroergonomics and cognitive computing* (pp. 315-326). New York: Springer.
- Cassimatis, N. (2012). Human-level artificial intelligence must be an extraordinary science. *Advances in Cognitive Systems*, 1, 37-45.
- Chen, S., & Epps, J. (2014). Using task-induced pupil diameter and blink rate to infer cognitive load. *Human-Computer Interaction*, 29, 390-413.
- Chomsky, N. (1959). A review of BF Skinner's *Verbal behavior*. *Language*, 35, 26-58.
- Cummings, M. M. (2014). Man versus machine or man + machine? *IEEE Intelligent Systems*, 29, 62-69.
- de Winter, J. C. F., & Hancock, P. A. (2015). Reflections on the 1951 Fitts list [sic]: Do humans believe now that machines surpass them? *Procedia Manufacturing*, 3, 5334-5341.
- Fahlman, S.E. (2012). Beyond idiot-savant AI. *Advances in Cognitive Systems*, 1, 15-21.
- Fitts, P. M. (1951). *Human engineering for an effective air navigation and traffic control system*. Washington, DC: National Research Council.
- Forbus, K. (2012). How minds will be built. *Advances in Cognitive Systems*, 1, 47-58.
- Gardner, H. (1987). *The mind's new science: A history of the cognitive revolution*. New York: Basic Books.
- Gigerenzer, G., Gaissmaier, W., Kurz-Milcke, E., Schwartz, L., & Woloshin, S., (2007). Helping doctors and patients make sense of health statistics. *Psychological Science in the Public Interest*, 8, 53-96.
- Jones, R. M., Wray, R. E., and van Lent, M. (2012). Practical evaluation of Integrated Cognitive Systems. *Advances in Cognitive Systems*, 1, 83-92.
- Jøsang, A. (2016). *Subjective Logic: A Formalism for Reasoning Under Uncertainty*. Cham, Switzerland: Springer.
- Kaber, D. B., Wright, M. C., Prinzel III, L. J., & Clamann, M. P. (2005). Adaptive automation of human-machine system information-processing functions. *Human Factors*, 47, 730-741.

- Cassenti, D. N., Veksler, V. D., Ritter, F. E. (2022). Editor's Review and Introduction: Cognition inspired artificial intelligence. *Topics in Cognitive Science*, 14, 652-664.
- Kahneman, D. & Tversky, P. (1973). On the psychology of prediction. *Psychological Review*, 80, 237-251.
- Kelley, T.D., & Cassenti, D. N. (2011). Theoretical explorations of cognitive robotics using developmental psychology. *New Ideas in Psychology*, 29, 228-234.
- Laird, J. E. (2012). *The Soar cognitive architecture*. Cambridge, MA: MIT Press.
- Langley, P. (2012). The cognitive systems paradigm. *Advances in Cognitive Systems*, 1, 3-13.
- Leito, A., & Radicioni, D. P. (Eds.). (2016a). From human to artificial cognition (and back): New perspectives of cognitively inspired AI systems. *Cognitive System Research*, 39, 1-72.
- Leito, A., & Radicioni, D. P. (2016b). From human to artificial cognition (and back): New perspectives of cognitively inspired AI systems. *Cognitive System Research*, 39, 1-3.
- Lobo, L., Heras-Escribano, M., & Travieso, D. (2018). The history and philosophy of ecological psychology. *Frontiers in Psychology*, 9, 2228.
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A proposal for the Dartmouth Summer Research Project on artificial intelligence, August 31, 1955. *AI Magazine*, 27, 12-14.
- McClay, W. J. (1995). Surviving the AI Winter. In J. Llyod (Ed.) *Proceedings of the International Symposium on Logic Programming* (pp. 33-47). Cambridge, MA: MIT Press.
- McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986). The appeal of parallel distributed processing. In D. A. Balota & E. J. (Eds.) *Cognitive Psychology* (pp 3-44). Cambridge MA: MIT Press.
- Miller, G. A. (1956). The magic number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 91-97.
- Moore, G.E. (1965). Cramming more components onto integrated circuits. *Electronics*, 38, 114.
- Newell, A. (1970). Remarks on the relationship between artificial intelligence and cognitive psychology. In R. Banerji & D. Mesarovic (Eds.). *Theoretical approaches to non-numerical problem solving* (pp. 363-400). Berlin: Springer-Verlag.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Peper, F. (2017). The end of Moore's Law: Opportunities for natural computing? *New Generation Computing*, 35, 253-269.
- Ritter, F. E., Tehranchi, F., & Oury, J. D. (2019). ACT-R: A cognitive architecture for modeling cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 10(3), Paper e1488.

- Cassenti, D. N., Veksler, V. D, Ritter, F. E. (2022). Editor's Review and Introduction: Cognition inspired artificial intelligence. *Topics in Cognitive Science*. 14. 652-664.
- Rotman, D. (2020) We're not prepared for the end of Moore's Law, *MIT Technology Review*. (Available at: <https://www.technologyreview.com/2020/02/24/905789/were-not-prepared-for-the-end-of-moores-law/>; Last accessed February 17, 2022)
- Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (Eds.). *Explorations in the Microstructure of Cognition, 1*, (pp. 45-76). Cambridge, MA: MIT Press.
- Sammur (2012). When do robots have to think? *Advances in Cognitive Systems, 1*, 73-81.
- Simon, H. A. (1977). Artificial intelligence systems that understand. In Proceedings of the *International Joint Conferences on Artificial Intelligence* (pp. 1059-1073). Cambridge, MA: MIT Press.
- Simon, H. A. (1981). Studying human intelligence by creating artificial intelligence: When considered as a physical symbol system, the human brain can be fruitfully studied by computer simulation of its processes. *American Scientist*, 69, 300-309.
- Tang, J., Yuan, F., Shen, X., Wang, Z., Rao, M., He, Y., Sun, Y. Li, X., Gao, B., Qian, H., Bi, G., Song, S., Yang, J. J, & Wu, H. (2019). Bridging biological and artificial neural networks with emerging neuromorphic devices: Fundamentals, progress, and challenges. *Advanced Materials*, 31, 1902761.
- Theraulaz, G., & Bonabeau, E. (1995). Coordination in distributed building. *Science*, 269, 686-688.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433-460.
- Zelazo, P. D. (2000). Self-reflection and the development of consciously controlled processing. In P. M. K. Riggs (Ed.), *Children's reasoning and the mind* (pp. 169–189). London: Psychology Press.
- Zelazo, P. D. (2004). The development of conscious control in childhood. *Trends in Cognitive Sciences*, 8, 12–17.

9. Acknowledgements

Marc Burns and Lance Kaplan provided comments that helped improve this paper. A special thanks to all of our colleagues who contributed articles to this special issue.