Advances in Cognitive Systems 5 (2017) 13–24 Submitted 11/2017; published 12/2017

Lessons from Human Problem Solving for Cognitive Systems Research

Alexandra Kirsch ALEXANDRA.KIRSCH@UNI-TUEBINGEN.DE
Department of Computer Science, Eberhard Karls Universität Tübingen, Tübingen, Germany

Abstract

Studies of human problem solving were one of the basic sources of insight for early AI research. Unfortunately, this approach has been largely abandoned in recent decades. In this essay, I argue that the cognitive systems community should put more effort into understanding and modeling human problem solving. I present examples of open research questions, discuss obstacles in pursuing this type of research, and show how greater diversity of methods and problems would improve understanding of on problem solving in both humans and machines.

1. Introduction

Attempts to understand human problem solving played a key role in early AI research. Seminal work of Newell and Simon (1972) on this topic paved the way for one of AI's most basic mechanisms: heuristic search through a problem space. Analyzing and drawing lessons from human behavior is more important than ever for research in AI and cognitive systems (Langley, 2012b; Lake et al., 2016). One reason is that people are the only entities we know that exhibit general intelligence. Psychological studies have revealed that abstract thinking is central not only to high-level cognition, but also to our ability to interpret events and react flexibly to them. The cognitive systems paradigm, which aims to revive the original vision of AI, adopts this perspective. If our goal is to implement computational systems that exhibit human-level intelligence, what is more obvious than looking at people's behavior for inspiration and guidance? Langley (2012a) explicitly includes links to human cognition as a characteristic of many cognitive systems efforts, although not all research in the area must make such connections.

Another motivation for understanding human behavior is the growing pervasiveness of embedded computers and, with it, the growing need for intuitive interfaces. Early computers were fixed to one location and people assumed that using them required special attention; in contrast, many modern computers are designed for use by anyone in any place. We can no longer count on highly-trained users who think like machines; we must devise machines that work like people, at least on the level of their visible behavior. AI's current focus on statistical machine learning has often led to situations in which not even the AI experts who build a system can understand its behavior, which in turn has led researchers to develop techniques for generating post-hoc explanations of opaque statistical models. In contrast, methods that are inspired by insights into the nature of human problem solving are more likely, from the outset, to be understandable without such convoluted and unnecessary rationalizations.

The gradual shift to statistical approaches over the last few decades has been accompanied by a change in the meaning of "problem". Newell and Simon (1972) studied tasks that were associated with a high degree of intelligence in humans (chess, cryptarithmetic puzzles, and logic). Unfortunately, such analytical problems are relatively easy to model mathematically and some AI researchers have recast them as optimization tasks. The result is that many people now associate the term "problem" with well-defined analytical tasks that have a well-defined optimum. Simon (1977) argues that most problems are inherently ill structured, even though their subtasks may involve well-defined structures. The original intention was to understand the solution of intellectually challenging problems in context by examining how humans approach them. In this view, problems are suitable for our purposes if humans are good at solving them.

In this essay, I use the Traveling Salesperson Problem (TSP) as an example. This task involves finding the shortest tour through a set of locations. It is both a classical problem and an abstraction of many real-world tasks, such as planning a vacation tour (Tenbrink & Seifert, 2011) and running errands in town (Hayes-Roth & Hayes-Roth, 1979), and it has been studied by cognitive psychologists (MacGregor & Chu, 2011; Best & Simon, 2000). First I present some examples of open questions in the area of human problem solving, showing that much remains to be done that has interesting scientific potential. Next I explain why this type of research is often difficult to pursue in the current academic environment. Finally, I argue that adopting the principles of design thinking may offer a way out of the current impasse. The core message is that science should foster diversity in research questions, paradigms, and methods.

2. Research Topics

Recent decades have seen AI producing remarkable results in well-defined but niche areas. This success led both to commercial exploitation of the responsible methods and increased research funding. As Fahlman (2012, p. 6) puts it, "When one of these super-human technologies takes off, it creates a 'gold rush' that attracts talent and resources away from the broader core problems of AI". Moreover, successes have been due partly to developments outside of AI, such as increased processor speeds, larger and cheaper memory, and introduction of infrastructure like the internet. This does not make them any less useful, but it suggests more rapid AI progress than has actually occurred. Many basic questions that have been open for decades are still not remotely solved.

In this section, I review some research questions that Newell and Simon (1972) raised 45 years ago, but that have received hardly any attention. I then propose some additional phenomena that are worth examining. I should emphasize that these are only examples; impact will not come from progress on them in isolation, but rather from advances on the entire set of problems and phenomena.

2.1 Uncharted Search Parameters

Newell and Simon (1972) identified search through a problem space as the basic mechanism of human problem solving. This paradigm is well accepted in AI, but many "parameters" allowed by the general framework have been forgotten. Russell and Norvig's (2010, p. 77) textbook presents an alluring abstraction in which search emerges entirely from the queuing function for expanded nodes. However, there are parameters that this scheme does not consider, which I will now examine.

Selective backtracking. There is no explicit notion of backtracking; search involves expanding a node that has been generated earlier, with the one selected being determined by a priority queue. Humans cannot use this strategy on even simple tasks because of their restricted short-term memory capacity. Think-aloud protocols (Newell & Simon, 1972; Hayes-Roth & Hayes-Roth, 1979) show that people sometimes reconsider earlier options, but only ones that "make sense" to them. "From time to time, subjects abandon the current information state they have reached and return to a prior state. They do not, however, retain the information that would permit them to return to *any* node they have visited previously. On the contrary, at any given point in the search only one or two nodes are commonly available as backup to the current one."(Newell & Simon, 1972, p. 815). For TSP solving, this translates into saving the unchosen action[1] when two alternatives are regarded as equally desirable. If later the path is found to be undesirable (e.g., when the only options that remain have crossing edges), they return to the unselected branch and try to find a solution that includes it. Figure 1 shows an example of how far people backtrack in this setting.

Branching. Russell and Norvig's (2010) technique expands the current node by applying every available operator. In an everyday context, this would mean that, when I think about how to walk to the nearest door, I would have to consider every movement of which my body is physically capable. Instead, I consider only a few well-trained leg movements. One can either put effort into choosing which operators to apply or into evaluation of resulting states. Some cognitive architectures (e.g., Prodigy, Soar, Icarus) explicitly consider the choice of operator. The idea is much less common in robotics, although a few researchers address it half-heartedly by sampling or filtering the actions with some criterion, as in the Dynamic Window Approach for navigation (Fox et al., 1997). But the explicit consideration of alternative operators is key to understanding how humans can solve NP-hard problems such as the TSP with relatively little effort, making it a more promising approach for incorporation in cognitive systems.

Abstraction. Newell and Simon (1972) observed that people sometimes switch between levels of abstraction, but they offered no final answer about how this occurs. Obviously, humans somehow abstract from a situation and ignore all but the "important" aspects. This process is often modeled by grouping states into clusters, sometimes known as "state abstraction", or organizing actions into higher-level activities, sometimes called "temporal abstraction", as in work on three-layer architectures, hierarchical task networks, and hierarchical reinforcement learning. Researchers have applied hierarchical approaches to the the TSP (Kong & Schunn, 2007; Best, 2006; Graham et al., 2000), but analysis of human wayfinding instead points towards an interplay of different abstraction levels (Wiener & Mallot, 2003; Tenbrink & Seifert, 2011). My own work (Kirsch, 2012) proposed a heuristic approach in which abstract knowledge, consisting of regional clusters, serves as only one factor for generating and evaluating alternative actions. This flat decision process used knowledge about spatially meaningful regions to the same advantage as a hierarchical method, but it was more stable when points were assigned randomly to regions. These examples serve to clarify that even well-established paradigms like heuristic search have not been fully exploited to explain and model the varieties of intelligent behavior.

1. Humans sove TSPs by incorporating one unvisited point after another into the tour, each step constituting an action.

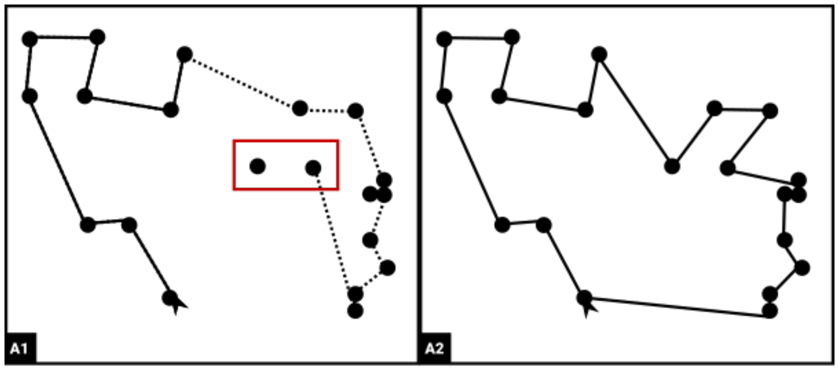

Figure 1. Trial of human behavior on the Traveling Salesperson Problem (Rach, 2017). The participant used the "undo" functionality eight times to revise an earlier decision (left), so as to include the two points marked in the red box earlier in the tour (right). This shows that backtracking is selective: the participant did not just go back one step for revision, but rather to a particular node in the search tree that seemed promising.

2.2 Other Problems and Phenomena

Search is not the only paradigm that needs attention and the TSP is neither the ideal or the only problem worth studying. But what are good problems to consider? And should we really focus on problems, or rather on capacities such as attention and decision making? There is no ideal problem that would give us enough insight to solve all the questions around cognition, or at least we are not advanced enough to identify such a problem, even if it existed. But while we can proceed only one step at a time, it is important to realize that individual problems are not usually the unit of interest. They are only instances that can provide insight into underlying capabilities. Therefore, I suggest that the research community should pursue a mixture of abilities on a mixture of problems.

One well-established phenomenon worth studying is the use of heuristics in decision making, following up on Newell and Simon's (1972) original research and revisiting common assumptions about the exploration of problem spaces. Search is a powerful concept, but it is well known that people only consider a small portion of the space. It is quite possible that human search is more like sampling, where representation of the space becomes more important and memory retrieval becomes a central factor (Jones & Langley, 2005). Note that this view differs drastically from mainstream AI's systematic exploration of search spaces. Recent psychological studies of heuristic decision making have occurred in behavioral economics (Samson, 2014). Although this field has revisited some basic assumptions about human behavior, it still lacks a theory of heuristics that is detailed enough for use in formal models.[2] Experimental designs typically involve single decisions with a specified set of alternatives. Computational models can enrich this research paradigm by proposing specific mechanisms, incorporating well-defined heuristics, and studying their effects.

Another alternative paradigm to search, related to heuristics, involves the concept of *habits*. In 2016, I co-organized an interdisciplinary workshop on this topic.[3] The meeting revealed that,

2. Gigerenzer (2001) and Shah and Oppenheimer (2008) present some promising accounts of heuristic decision making.
3. The workshop on "Technology and Routines at the Individual and Organizational Level", held at the Bavarian Academy of Sciences and Humanities, was co-organized with Jutta Stumpf-Wollersheim.

although the field of economics distinguishes between (individual) habits and (organizational) routines, I could not find any recent AI publication that explicitly mentioned these ideas. Miller et al. (1960) explored the topic in the seminal book, and one could argue that case-based reasoning uses habits, but current research in the area does not make this connection (Richter & Weber, 2013). Since habits are such an integral part of human cognition, we should try to understand them as a powerful mechanism that reduces cognitive and computational load. Their use can also produce more predictable behavior and let systems classify, understand, and adapt to human regularities.

Habits depend on representing and storing sequences of actions. AI offers frameworks such as HTNs, temporal logic, and Markov chains, but such sequences involve many attributes that must be encoded, including the order of actions, their duration, their absolute times, and their locations (Roor et al., 2017b). One complication in analyzing activities is that humans follow habits loosely; another is that they divide sequences into meaningful subsequences, such as the routine of preparing for work and another for going to one's workplace. Psychologists have found interesting patterns in how people parse sequences in movies (Huff et al., 2014) that could inspire computational systems. Such phenomena arise in different contexts. For instance, habits are the basis of many everyday tasks and can be observed in smart environments, and mobile devices offer a plethora of sensor information that could be used to study habitual behavior (Roor et al., 2017a).

In addition to such complex, real-world tasks, researchers can study targeted problems that lend themselves to greater experimental control. Games have long been a favorite subject for research, as they combine a set of well-specified rules, mostly observable environments, and well-defined goals. In addition, people typically enjoy playing games, which makes it reasonably easy to collect data. We have implemented a game that makes users solve instances of the TSP (Rach & Kirsch, 2016) for precisely this reason.

3. Obstacles to Progress

If there are so many open questions about problem solving, why is more research not being done in this area? In this section, I describe three obstacles to progress: the difficulty of finding appropriate problems to pursue, which holds for all of AI and cognitive systems research; the narrow computer science view on solution quality in terms of formal optimization criteria; and the research environment, in which scientific quality is equated with publishing and fund-raising success. These obstacles reduce the number of career opportunities and limit the funding available to researchers who develop cognitive systems compared to those who make incremental enhancements to existing algorithms for narrow tasks. In Section 4, I propose one way out of this impasse.

3.1 The AI Dilemma

In AI and cognitive systems, one of the most difficult challenges is to find appropriate problems to tackle. If a problem is too simple, then it can always be engineered and thus does not really require AI methods. If it is too complex, then its solution requires a variety of techniques and their appropriate combination; thus, success or failure cannot be attributed to any single component.

This dilemma applies especially to understanding problem solving: if we examine problems from a narrow angle, we can model anything. The result of a psychological experiment can be directly turned into a computational model that mimics exactly the observed behavior. Unfortunately, such narrow models fail to explain intelligence. Developing more general accounts encounters at least two basic hurdles. First, psychologists take great pains to control conditions in their experiments. This makes absolute sense in order to identify specific situations that produce specific behaviors. But it is a small, well-defined setting, and any small, well-defined problem can be engineered; this does not require us to understand intelligence. Second, showing that a computational approach is general requires its application to, and testing on, different problems. Since the whole purpose of such an approach is not to "overfit" to a specific task, it will typically perform worse on any particular problem than methods that are handcrafted for that problem.

I can illustrate this point with the TSP example, which I have been trying to model with a heuristic decision-making approach for several years (Kirsch, 2011, 2012; Rach & Kirsch, 2016). On one hand, it is too simple, in that it involves a fully observable, abstract problem that, in its pure form with points and lines, has little to do with everyday human tasks. And it can be solved effectively by algorithms that find optimal solutions for problems with thousands of points and near-optimal answers for ones with millions of points, which is far beyond human capabilities.

On the other hand, the TSP is too complex. I chose this problem because I wanted to study decision making with minimal concern for issues of knowledge representation. But I realized that, when people solve (artificially presented) TSPs, they seem to treat it as a pattern-matching exercise, trying mainly to produce an aesthetic figure out of the points rather than optimizing for a short tour (Vickers et al., 2006). To model more realistic variants of the task, one would need to consider more cognitive processes and forms of knowledge representation. Do we use the same mechanisms to solve abstract TSPs as when planning a vacation? How do people represent distances, points of interest, and preferences? Also, evaluation is a challenge. Even if we could clearly specify whether a specific TSP instance is solved well enough, it would not tell us much about the heuristic solution process. The success or failure of solving TSPs could be due to the heuristics themselves or to the decision mechanism that uses them.

3.2 The Myth of Optimality

I have broadened the definition of "problem" to include any task that arises in the context of real-world activities. This is a rather big step for many AI researchers, who are used to dealing with well-defined tasks that may be intractable in computational terms but at least have a well-defined solution or optimum. Simon (1993) explains that there are three steps in decision making: the choice of problems or subtasks to attend to; the generation of alternatives; and the selection of an alternative. Like economists, AI researchers usually focus on the third step and often define rationality as choosing the alternative with the highest expected utility. However, "[m]aximizing utility bears no resemblance whatsoever to what we human beings actually do".

A critical factor is the knowledge involved. Gigerenzer and Gaissmaier (2011, p. 452–453) refer to Savage (1954) to distinguish between 'small' and 'large' worlds: "In large worlds, part of the relevant information is unknown or has to be estimated from small samples, so that the conditions for rational decision theory are not met, making it an inappropriate norm for optimal reasoning (Binmore, 2009)". This differentiation may at first glance be just a technical one. But the critical point is that most researchers in AI implicitly assume that methods which are optimal in small worlds also produce good solutions in large worlds. Gigerenzer and Gaissmaier provide compelling examples that "small-world theories can lead to disaster when applied to the large world".

A common fallacy is to ignore the modeling step and then expect optimal solutions. Whenever engineers must solve a problem in the real world, they first build a more or less detailed model of the task. For example, one could model the planning of a vacation trip as a TSP, maybe with a modified cost function to account for personal preferences or beauty of the landscape. From this point on, most AI researchers would be obsessed with finding optimal solutions. Suboptimal solutions are only accepted due to intractability. However, a model never represents the world perfectly and there is no way to establish a formal mapping between reality and the model. If there is such a subjective step in the process, why should we insist on optimality for the rest of the process?

One answer is that evaluation of a method is much easier for the model than for the original task. Verifying whether a method provides good answers in everyday tasks requires feedback from users and measurements in the environment. Thus, we would have to move from a clean mathematical task into messy empirical research. Again consider the Traveling Salesperson Problem as an example. In the model world, the quality criterion is clear: the length of the tour. But most real-world TSP variations cannot be mapped onto such a simple measure. If computers are to give advice to people, we should have some measure of comparing human solutions with computational ones. Tak et al. (2008) point out that, even for classical TSPs, length of the tour is an inadequate measure, because tours with the same length can differ in structure and vice versa. Also, number of crossings has limited use, because humans rarely produce solutions with crossing edges. The next best thing would be to find different ways to compare tours structurally, say by the number of edges that overlap, but any such method is rather arbitrary. The ultimate test for a TSP method would be its application to a generalized TSP task in which costs are less clear than in the formalized version, then to have people rate the quality of solutions. But this would require the effort of running a psychological experiment and reduce the result to the specific instance tested in the study.

Note that this type of evaluation disadvantages, in three ways, researchers who are not willing to adopt a mathematical model with well-defined optimality criteria. First, they need more time to do their research, reducing the number of publications produced per time unit. Second, methodological flaws are more likely to occur in a paper covering the whole pipeline, from modeling a problem to testing an implemented approach empirically, than in one that solves a well-defined problem theoretically under unrealistic assumptions. This again makes it more likely that a paper will be rejected. Third, if the ultimate test lies in observing human responses, then the testing criterion is not accessible at the time of method development. Thus, such a test may produce "negative" results, in the sense that the method fares worse than expected or produces worse results than existing, more finely-tuned niche solutions. From a scientific perspective, such results are just as valuable as ones finding that a method works well, but they are more difficult to publish than reports of incremental improvements on specific tasks.

3.3 The Research Environment

I have shown how the unique dilemma faced by AI requires the study of a diverse spectrum of problems and phenomena. Overcoming the myth of optimality would require a variety of methods, particularly for evaluation. Unfortunately, in a world where researchers are pressured to publish as much as they can in the smallest time possible, this diversity is difficult to achieve. Many AI subfields agree on a specific set of methods, often with some notion of optimality. Other methods are regarded as unscientific, leading to rejection of papers and research proposals. Some fields even rely on a common set of problems. For example, in the AI planning literature, problems used for evaluation are often unexplained; they are simply references to the latest planning competition.

The preference for niche-AI research over broader topics has its analog in behaviorist approaches to psychology. Both are connected to a general trend in science that favors reductionist explanations, which has been criticized in multiple scientific fields (Breckler, 2006; Kaiser, 2011). Hopkins et al. (2016) have shown that reductionist explanations are often regarded as more scientific even when they are irrelevant to the problem. Thus, some types of explanations appear more scientific than others and receive preferential treatment.

Moreover, science seems to be echoing a trend in Western societies by following the narrative of ever-faster, ever-bigger success. To aggravate the situation, researchers face increased competition as research funding focuses more on short-term results than on long-term progress and basic research infrastructure. This is leading to a veritable arms race among researchers for who has the longest publication list and the most funding. This in turn encourages researchers to choose topics that can be published quickly and easily, and ones that fit well into the "mainstream" of their research paradigm (Lawrence, 2007).

One rare example where findings from human problem solving have been used is research on cognitive architectures. Some architectural theories have a development history of decades. In general, good research needs time to think thoroughly about possible implementations and to discover which methods work and which do not. Geman and Geman (2016) compare the current practice of science with constantly taking selfies: "In fact, many of us spend more time announcing ideas than formulating them. Being busy needs to be visible, and deep thinking is not".

In this context, research on cognitive systems and human problem solving entails a high risk of being disadvantaged in funding and career decisions: identifying suitable problems and performing empirical tests takes more time and is less highly regarded than work on well-defined, established problems with formal evaluation criteria. This leads directly to lower acceptance of papers and research proposals. However, if enough researchers agreed on the importance of this type of research, we could build a community with a different set of values and methods. For example, the online journal and annual conference *Advances in Cognitive Systems* provide venues for research that follows the original quest for general-purpose, human-level AI (Langley, 2012a).

4. Design Thinking for Cognitive Systems Research

Winograd (2006) draws a line between AI and human-computer interaction not so much in their research questions, but in terms of their methodologies. The HCI community uses ideas from design research to foster diversity of ideas. "Design thinking" has become a buzzword for an iterative development process in which divergent and convergent thinking alternate, with a strong emphasis on prototyping. This approach contrasts with following steps in a preconceived plan or script. In software development, this change has been instantiated by a shift away from linear development frameworks like the waterfall model towards agile software development.

A basic feature of design thinking is to place the user and task at the center of attention. As the implementation details of products is not a central concern of scientific research, this may seem unimportant. But many researchers in niche AI communities have forgotten or never even considered the larger context to which their research contributes. Design thinking can remind us to take time on occasion to consider why we are examining a certain intellectual ability or why we are trying to solve a certain class of problems.

As noted, the design process alternates between *divergent* and *convergent* thinking. Niche AI typically emphasize the latter in its efforts to improve methods incrementally. However, this should not be a blind optimization process, but rather a sequence of conscious decisions about which incremental steps lead to the most desirable improvements and whether the current method needs substantial changes or extensions. In addition to convergent thinking, the design process emphasizes divergent cognition: proposing many alternative ideas, either for the construction of complete intelligent systems or for automating component abilities. The design movement encourages people to think not only along known lines, but also to "think outside the box" and even consider unlikely or crazy ideas. The approach draws on prototyping to determine quickly and cheaply which alternatives are most promising for further exploration. This explicitly includes the possibility of following blind alleys, at least for limited time, and recognizes the inability to predict the consequences of any alternative without having tried them in enough detail. Therefore, determining that some technique does not work as expected is just as valuable as finding that it does.

A major challenge in this process is determining whether a path is promising and which directions to extend it. The only answer I can give, which may be unsatisfying to scientists familiar with quantitative measures and fixed evaluation procedures, is to rely on common sense. Science involves searching through a very large space and we need heuristics to find our way. Luckily, the human mind comes equipped with a set of very useful heuristics; these fail in certain cases (Tversky & Kahneman, 1974), but they usually lead to appropriate decisions (Gigerenzer & Brighton, 2009).

The cognitive architecture paradigm – a direct descendant of Newell and Simon's research on human problem solving – offers a good example of design thinking (Langley, 2017; Langley et al., 2009). Research in this area aims to develop unified theories of human cognition and/or intelligent agents, which it pursues over long periods (Kotseruba & Tsotsos, 2016). Convergent thinking may refine specific elements of an architecture, such as mechanisms for efficient memory retrieval. But divergent thinking reminds one tht each such component function can be instantiated in different ways, such as sorting items through a discrimnation network or spreading activation through memory. Moreover, these elements are subject to review and change, with entire architectures being tested on a wide range of problems that serve to ensure their generality.

Cognitive architectures provide a good example of a research paradigm that satisfies many of my criteria, but we also need diversity at more detailed levels. The holistic view of cognitive architectures should be complemented by reductionist approaches, to better understand how specific mechanisms can be integrated into a unified theory once they have proved promising. However, in the current research environment, reductionist schemes are overemphasized and pursued outside such a holistic context. Although not everyone must develop a cognitive architecture, they should be aware that component techniques are building blocks for complete systems, and they should be developed and evaluated in this broader context.

5. Conclusion

In this essay, I argued that the connection between research on human problem solving and artificial intelligence has been largely abandoned, while theoretical optimization has gained importance. This has led to a situation in which holistic progress is hardly visible from outside niche fields, while other important topics receive little or no attention. To some extent, the cause lies in the difficulty of AI itself, especially the key task of finding appropriate problems. This inherent challenge has been aggravated by goals and incentives that misdirect AI research and science more generally.

This impasse can only be overcome by pursuing greater diversity of both methods and problems. This requires an openness of mind from researchers, but it can also be supported by the paradigm of design thinking. With a clear focus on the original goals of AI and cognitive systems, every type of research, be it a single method or a complete cognitive architecture, should be regarded in the context of the overall endeavor. A mixture of divergent and convergent thinking, with conscious decisions about how to proceed, would lead to increased diversity and progress in the field.

When researchers test their methods on different problem instances, such tests will sometimes reveal that a method does not work well. The concept of prototyping captures the benefits of trial and error; if more researchers embraced it, concerns about "negative" results would arise less often or disappear altogether; any result would be equally welcome, since it advances the field's knowledge. In short, embracing prototyping by applying different research methods to various problems would lead to more robust systems and help banish the notion of negative results.

Acknowledgements

I would like to thank Pat Langley and Don Rose for reading and commenting on earlier drafts of this essay. The research reported in this paper was supported by the Bavarian Academy of Sciences and Humanities and the German Research Foundation.

References

Best, B. J. (2006). Cognitive approaches to the traveling salesperson problem: Perceptual complexity that produces computational simplicity. *Papers from the AAAI Spring Symposium on Cognitive Science Principles Meet AI-Hard Problems* (pp. 11–16). Stanford, CA: AAAI Press.

Best, B. J., & Simon, H. A. (2000). Simulating human performance on the traveling salesman problem. *Proceedings of the Third International Conference on Cognitive Modeling* (pp. 42–49). Groningen, Netherlands: Universal Press.

Binmore, K. (2009). *Rational decisions*. Princeton, NJ: Princeton University Press.

Breckler, S. J. (2006). The newest age of reductionism. *Monitor on Psychology*, *37*, 23.

Fahlman, S. E. (2012). Beyond idiot-savant AI. *Advances in Cognitive Systems*, *1*, 15–21.

Fox, D., Burgard, W., & Thrun, S. (1997). The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine*, *4*, 23–33.

Geman, D., & Geman, S. (2016). Science in the age of selfies. *Proceedings of the National Academy of Sciences*, *113*, 9384–9387.

Gigerenzer, G. (2001). The adaptive toolbox. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox*, 37–50. Cambridge, MA: MIT Press.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, *1*, 107–143.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451–482.

Graham, S. M., Joshi, A., & Pizlo, Z. (2000). The traveling salesman problem: A hierarchical model. *Memory & Cognition*, *28*, 1191–1204.

Hayes-Roth, B., & Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science*, *3*, 275–310.

Hopkins, E. J., Weisberg, D. S., & Taylor, J. C. (2016). The seductive allure is a reductive allure: People prefer scientific explanations that contain logically irrelevant reductive information. *Cognition*, *155*, 67–76.

Huff, M., Meitz, T. G., & Papenmeier, F. (2014). Changes in situation models modulate processes of event perception in audio-visual narratives. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*, 1377–1388.

Jones, R. M., & Langley, P. (2005). A constrained architecture for learning and problem solving. *Computational Intelligence*, *21*, 480–502.

Kaiser, M. I. (2011). The limits of reductionism in the life sciences. *History and Philosophy of the Life Sciences*, *33*, 453–476.

Kirsch, A. (2011). Humanlike problem solving in the context of the Traveling Salesperson Problem. *Papers from the AAAI Fall Symposium on Advances in Cognitive Systems* (pp. 146–153). Stanford, CA: AAAI Press.

Kirsch, A. (2012). Hierarchical knowledge for heuristic problem solving: A case study on the Traveling Salesperson Problem. *Poster Collection: The First Annual Conference on Advances in Cognitive Systems* (pp. 87–102). Palo Alto, CA: Cognitive Systems Foundation.

Kong, X., & Schunn, C. D. (2007). Global vs. local information processing in visual/spatial problem solving: The case of traveling salesman problem. *Cognitive Systems Research*, *8*, 192–207.

Kotseruba, I., & Tsotsos, J. K. (2016). *A review of 40 years of cognitive architecture research: Core cognitive abilities and practical applications*. cs.AI arXiv:1610.08602v2.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2016). Building machines that learn and think like people. *Behavioral and Brain Sciences*, *24*, 1–101.

Langley, P. (2012a). The cognitive systems paradigm. *Advances in Cognitive Systems*, *1*, 3–13.

Langley, P. (2012b). Intelligent behavior in humans and machines. *Advances in Cognitive Systems*, *2*, 3–12.

Langley, P. (2017). Progress and challenges in research on cognitive architectures. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* (pp. 4870–4876). San Francisco: AAAI Press.

Langley, P., Laird, J. E., & Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, *10*, 141–160.

Lawrence, P. A. (2007). The mismeasurement of science. *Current Biology*, *17*, R583–R585.

MacGregor, J. N., & Chu, Y. (2011). Human performance on the traveling salesman and related problems: A review. *The Journal of Problem Solving*, *3*, 1–29.

Miller, G. A., Galanter, E., & Pribram, K. A. (1960). *Plans and the structure of behavior*. New York: Holt, Rinehart, & Winston.

Newell, A., & Simon, H. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Rach, T. (2017). *Heuristische modellierung menschlicher lösungsstrategien für das Traveling Salesperson Problem*. Master's thesis, Department of Computer Science, Eberhard Karls Universität, Tübingen, Germany.

Rach, T., & Kirsch, A. (2016). Modelling human problem solving with data from an online game. *Cognitive Processing*, *17*, 415–428.

Richter, M. M., & Weber, R. O. (2013). *Case-based reasoning: A textbook*. Heidelberg: Springer.

Roor, R., Hess, J., Saveriano, M., Karg, M., & Kirsch, A. (2017a). Sensor fusion for semantic place labeling. *Proceedings of the Third International Conference on Vehicle Technology and Intelligent Transport Systems* (pp. 121–131). Porto, Portugal: Springer.

Roor, R., Karg, M., Liao, A., Lei, W., & Kirsch, A. (2017b). Predictive ridesharing based on personal mobility patterns. *Proceedings of the Twenty-eighth IEEE Intelligent Vehicles Symposium* (pp. 1322–1329). Los Angeles: IEEE.

Russell, S. J., & Norvig, P. (2010). *Artificial intelligence: A modern approach* (3rd ed.). Upper Saddle River, NJ: Pearson Education.

Samson, A. (2014). An introduction to behavioral economics. In A. Samson (Ed.), *The behavioral economics guide 2014*, 1–12. London: Behavioral Science Solutions.

Savage, L. J. (1954). *The foundations of statistics* (2nd ed.). New York: Dover.

Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, *134*, 207–222.

Simon, H. A. (1977). The structure of ill-structured problems. In H. A. Simon (Ed.), *Models of discovery and other topics in the methods of science*, 304–325. Dordrecht: Springer Netherlands.

Simon, H. A. (1993). Decision making: Rational, nonrational, and irrational. *Educational Administration Quarterly*, *29*, 392–411.

Tak, S., Plaisier, M., & van Rooij, I. (2008). Some tours are more equal than others: The convex-hull model revisited with lessons for testing models of the Traveling Salesperson Problem. *The Journal of Problem Solving*, *2*, 4–28.

Tenbrink, T., & Seifert, I. (2011). Conceptual layers and strategies in tour planning. *Cognitive Processing*, *12*, 109–125.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.

Vickers, D., Lee, M. D., Dry, M., Hughes, P., & McMahon, J. A. (2006). The aesthetic appeal of minimal structures: Judging the attractiveness of solutions to traveling salesperson problems. *Perception & Psychophysics*, *68*, 32–42.

Wiener, J. M., & Mallot, H. A. (2003). 'Fine-to-Coarse' route planning and navigation in regionalized environments. *Spatial Cognition and Computation*, *3*, 331–358.

Winograd, T. (2006). Shifting viewpoints: Artificial intelligence and human-computer interaction. *Artificial Intelligence*, *170*, 1256–1258.