Advances in Cognitive Systems 5 (2017) 3–12 Submitted 12/2017; published 12/2017

Heuristics and Cognitive Systems

Pat Langley PATRICK.W.LANGLEY@GMAIL.COM
Institute for the Study of Learning and Expertise, Palo Alto, California 94306 USA
Department of Computer Science, University of Auckland, Private Bag 92019, Auckland 1142 NZ

Abstract

In this essay, I consider the important role of heuristics in research on artificial intelligence and cognitive systems. After clarifying the different senses of this term, I recount how views on heuristics have changed since their introduction, leading many in the AI community to see them in a very different light than intended originally by the field's founders. In addition, I present four claims about how heuristic methods and structures can influence high-level cognition and how these postulates differ from mainstream views on the topic. In closing, I propose some actions that cognitive systems researchers can take to redress the situation and restore heuristics, in the original meaning, to their rightful place in the computational study of intelligence.

1. Variations on a Meme

One of the central features of cognitive systems research is its incorporation of *heuristics*. This notion was central to early work on artificial intelligence and crucial to differentiating it from contemporary developments in computer science. In some cases, AI research was even equated with *heuristic programming* (e.g., Slagle, 1971), a phrase that contrasted the field with the algorithmic methods associated with other areas of computing. This suggests that the topic deserves closer attention from the primary inheritor of AI's original vision – the cognitive systems community.

In the sections that follow, I discuss the origins of this meme and state some theoretical claims about heuristics' relation to intelligent behavior. These will be familiar to many readers of this journal, as they are often assumed in cognitive systems research. But they are less widely adopted by what has become mainstream AI, so in each case I also examine how views on the topic have changed since the field's inception in the 1950s. Before proceeding further, I should clarify what I mean by the term *heuristic*. A careful reading of the AI and psychology literature reveals at least three distinct but related senses, and they deserve a brief review before I turn to their implications for AI and cognitive systems. I will not propose formal definitions in this essay, as they could distract from conceptual matters, but I will attempt to state things clearly enough to avoid ambiguities.

The most generic sense of *heuristic* refers to some method or strategy for making a decision or solving a problem that typically produces useful results at reasonable computational cost. Such techniques do not invariably find the best solution or, indeed, any solution at all, but they often work well in practice. These are often contrasted with *algorithms* that are guaranteed to find solutions, in some cases 'optimal' ones. A classic example from decision making is the *take the best* strategy (Gigerenzer & Goldstein, 1996) for selecting an item from a set of alternatives. Here one orders attributes by importance and chooses an item based on the most important attribute that discriminates among them. Another example involves *greedy search* (Cormen, Leiserson, & Rivest, 1990) in which one attempts to solve a multi-step problem by selecting the best alternative on each step. Reasoning by analogy (Gentner & Forbus, 1991; Polya, 1957) is another familiar heuristic method. Each approach often produces acceptable results while requiring low computational resources.

A narrower sense of the term *heuristic* is a rule of thumb. This is often associated with frameworks like production systems that encode knowledge as condition-action rules, although they may be stated in any symbolic formalism, from Prolog clauses to larger scripts or frames. Such structures can encode various forms of content that aid decision making and search. On simple choice tasks, rules that specify preferences (e.g., for solid colors over stripes) can be combined with a *take the best* strategy to decide among shirts. For problem-solving tasks like planning and design, rules that prefer, select, and reject goals, states, and operators can guide search through large spaces. Experiments with Soar (Laird, Newell, & Rosenbloom, 1987) and PRODIGY (Carbonell, Knoblock, & Minton, 1990) have demonstrated such control rules can be effective in solving difficult problems.

A third meaning of *heuristic* refers to evaluation criteria. These are central to decision-theoretic approaches to simple choice tasks, where they are called *utility functions*, but they also appear regularly in work on heuristic search and game playing. These play the same role as symbolic rules of thumb that specify preferences, but they produce quantitative scores for choices rather than a qualitative ranking on them. They also assume a very different format, typically an equation that combines a set of numeric features or attributes that are provided, measured, or calculated. Evaluation functions are favored by AI researchers who are attracted to continuous mathematics, and they have been used successfully not only for problems that involve sequential action, like planning and game playing, but are also commonly adopted for tasks like design and scheduling.

Each sense of the term is equally legitimate, and the meaning intended in a given publication is usually clear from its context. Nevertheless, it is still important to acknowledge the distinctions and to understand how they relate to each other, as this will help in stating theoretical claims about how heuristics relate to structures and processes that arise in cognitive systems. However, the importance of these postulates will be more transparent if I first review how the heuristic approach began and how views about it have changed over time.

2. The Development of Heuristic Thought

As already noted, heuristic processing was an early defining feature of AI research and played a valuable role in the field's development, but the idea originated in disciplines that predated the AI revolution of the late 1950s.[1] For example, the psychologist Selz (1927) proposed an early account of human problem solving in terms of tasks to be solved, mental transformations, and goal-oriented 'schematic anticipations'. As ter Hark (2010) has observed, these foreshadowed later ideas on heuristic search through a problem space. Similarly, Simon's (1947, 1956) concept of *satisficing*, which he introduced initially under a different name, also prefigured the later adoption of heuristic techniques. Both sets of ideas found their way into early AI systems like the Logic Theorist (Newell, Shaw, & Simon, 1957), although it is unclear whether the latter's authors were aware of Selz's work at development time (Simon, 1981).

1. Less formal usage of the term goes back much further to the original Greek word *heuriskein*, which means *to find*.

Another precursor was Polya (1957), whose well-known book, *How to Solve It*, contained advice for addressing problems in his field, mathematics. He used the word *heuristic* to refer to a diverse set of methods that, although not stated in fully operational terms, include many techniques of considerable generality. Although the book did not appear until after implementation of the first running AI system, Polya taught courses on the topic in the 1940s and Newell (1983) reports taking one when he was a freshman. Newell also notes that he was not consciously aware of this work when designing the Logic Theorist, but nevertheless believed that it influenced his thinking. Many AI researchers have read Polya's book, but very few have attempted to incorporate his detailed heuristic methods into their systems.

In many ways, early research on AI proceeded independently from efforts on what came to be known as *computer science*, with the first championing heuristic approaches and the second emphasizing algorithmic methods. Indeed, until the 1980s, both paradigms would have viewed the phrase *heuristic algorithm*, now widely accepted, to be an oxymoron. Remember that few departments of computer science were founded until the 1970s, which meant that many seminal dissertations on artificial intelligence were done with little exposure to views later associated with computer science and vice versa. Only in the 1980s, when computer science departments grew rapidly and began to hire AI researchers, regardless of their background, did the two communities come into intimate contact. This was partly a marriage of convenience, as computer science needed energetic professors and AI researchers had few other options for academic homes.

Unfortunately, the resulting union had negative impacts on some of AI's central tenets, especially its commitment to heuristic solutions. Many departments of computer science had grown out of mathematics units, which led formal analysis and provable results to receive higher priority than alternatives. Once the majority of the AI community had been trained in such settings, many adopted the view that, unless an approach offered formal guarantees, it was unprincipled and ad hoc. In such circles, heuristic approaches have become *persona non grata*, despite their crucial role in AI's early development and their clear advantages in building intelligent systems. This climate indicates a need to state explicitly some theoretical postulates about heuristics that, 30 years ago, would have been obvious, but that have been rejected or forgotten by modern researchers.

3. Heuristics and Tractability

As already implied, a key characteristic of heuristic methods relates to their functional role in cognition. High-level mental processing – the types of intellectual activities that separate humans from other mammals – often arises in complex task settings. Indeed, formal analyses in the subfield of complexity theory have shown that some problems cannot, in the worst case, be solved in reasonable times. However, in many situations, humans find these same problems quite manageable (e.g., Kirsch, 2011, 2017), despite having very limited information-processing resources.

This observation leads directly to our first claim, which attempts to explain people's ability to solve apparently difficult tasks and which many readers will find familiar:

- *Heuristic methods often make tractable the solution of apparently complex tasks.*

This postulate was a key motivation for heuristics' early association with the AI revolution, which dared to study problems that contemporary fields, such as computer science and operations research, assumed were too difficult for automation. All three disciplines now acknowledge the power of heuristic approaches, although we will see that they often adopt a narrower sense than intended by the founders of artificial intelligence.

Some readers may think the above claim is a tautology, as its phrasing is similar to one of the definitions given earlier. Certainly a common *reason* for invoking heuristic methods is to let one handle a complex problem with manageable effort, but they may not achieve this aim. Some strategies may be more effective than others, and some may work well only on certain classes of problems. Moreover, one can always take a useful strategy and 'invert' its choices to produce a highly ineffective one, meaning that 'bad' techniques are also possible. The implication is that whether a given heuristic method works in practice is an empirical question rather than a mathematical one, but many studies have produced results that support the conjecture.

Despite this evidence, the AI community has exhibited a growing bias toward formal results that has led many researchers to denigrate heuristic approaches. One result has been a shift toward restricted representations, such as description logics, that come with assurances of efficient computation (e.g., Levesque & Brachmann, 1987). Another has been an increased emphasis on classes of problems that one can solve with methods like dynamic programming (e.g., Kaelbling, Littman, & Cassandra, 1998) and even exhaustive search, which offer guarantees of optimal solutions. Publications in these paradigms often dismiss heuristics as 'ad hoc' and undesirable, revealing that they have abandoned one of the original tenets of artificial intelligence.

Similar attitudes have led to a subtler change in some subfields' notions of heuristics that, in certain ways, is still more disturbing. This view had its origins in Hart, Nilsson, and Raphael's (1968) invention of *A**, one of the first nonexhaustive search methods guaranteed to find solutions that were optimal in the sense of having the fewest steps or lowest cost. However, this alluring result holds only when the technique uses *admissable* heuristics, a form of evaluation function that underestimates the distance or cost to the goal or target state. Note that this name, in itself, suggests that all *other* heuristics are somehow inadequate and irrelevant. This seems strange given that A* provides no guarantees about reduction in search, which was AI's original reason for championing heuristics, and in some cases it can take longer than breadth-first methods.

Nevertheless, the A* algorithm is often presented in courses as the only reasonable approach to heuristic search, and it has come to dominate research in the AI planning and search communities. Some authors now treat the term *heuristic* as synonymous with *admissable heuristic*, suggesting that any other types are not worth mentioning. This usage is a true perversion of the word's original sense, which referred to methods that offered no guarantees but typically worked better in practice than ones which provided them. Thus, the introduction of A* was the beginning of the end for classical AI, and its widespread, uncritical adoption is a sign of how far the field has fallen from its early vision and tenets. This unfortunate change is one of the primary reasons why the cognitive systems paradigm now requires distinct publication venues from its ancestral field.

4. Heuristics and Satisficing

I have argued that heuristic methods can offer substantial advantages over 'algorithmic' ones, letting an AI system that draws on them solve quite complex problems with reasonably little effort. There is also overwhelming evidence that humans employ such methods when confronted with problem-solving tasks, such as playing chess (de Groot, 1978) and solving puzzles (Newell & Simon, 1972), and even on simple choice tasks, such as deciding which products to buy in a grocery store (Weber & Johnson, 2009). This offers another reason for studying heuristic approaches to complex mental behaviors and suggests a second claim:

- *Heuristic methods produce human-like cognition, especially satisficing behavior*.

This latter concept is worth elaborating, as many AI researchers have either abandoned it unconsciously or rejected it explicitly. Simon (1956) introduced the idea of *satisficing* in response to economists' widespread assumption that people make optimal choices. His studies revealed that, in all but the most trivial cases, they instead settle on finding an alternative that is *good enough*. For example, when buying a house or finding a job, we do not consider all options in detail, but rather examine a manageable subset, and we halt when we find one that is satisfactory. Simon linked this satisficing behavior to the notion of *aspiration level*, which informs a decision maker when a candidate is sufficiently high quality to halt processing. Again, many empirical studies of humans have revealed this behavior in a wide range of settings.

Naturally, this does not mean that heuristic methods can *never* find the most desirable solutions. Gigerenzer and Goldstein (1996) have demonstrated that simple strategies like *take the best* often make the same choices as decision-theoretic techniques that consider every alternative, calculate numeric values for them, and select the one with the highest score. What heuristics do not offer are *guarantees* that they will make the optimal choice. This causes great concern for AI researchers who have 'theorem envy' of mathematicians and complexity analysts, but it should not be an issue for the cognitive systems community, which focuses on human-like behavior and on mechanisms that work well empirically. In rare cases, heuristic methods may not find *any* solution, much less the best one, but this also holds for human problem solvers, who manage to muddle along anyway. Anyone who insists that an AI system is not rational unless it makes optimal choices is taking a dubious position, as this implies that people – our only proof that intelligence is possible – are not themselves rational.[2] Newell's (1982) definition – that *an agent is rational if it selects actions that it believes will lead to one of its goals* – seems a far more appropriate stance.

Ironically, heuristic methods and satisficing have even been undermined by cognitive psychologists. For instance, influential work by Tversky and Kahneman (1974) acknowledges the important role of heuristics in human decision making, including the influence of a choice's availability, representativeness, and other features on whether a person selects it. However, as Gigerenzer (1996) has noted, they view such heuristics as *biases* that cause divergence from the 'correct' choice, which reveals a bias toward decision-theoretic frameworks. Another example is Kahneman and Tversky's (1979) *prospect theory*, which differs from normative approaches like expected utility theory by introducing an aversion to loss. Yet their account still assumes that decision makers inspect all attributes and calculate weighted sums when making choices, contrary to strong evidence that people use simpler strategies which require less mental effort (Simon, 1956; Gigerenzer & Goldstein, 1996; Weber & Johnson, 2009). In contrast, the cognitive systems community offers a venue for research that treats humans, our one example of general intelligence, as a worthy target.

2. A common response is that one can extend the definition of optimality to incorporate the computational costs needed to make a decision, but this argument do not hold water. Humans certainly can take computational factors into account, but they use heuristic estimates in this meta-level context as well.

5. Heuristics and Knowledge

I have explained that one sense of *heuristic* is an adjective which describes methods or mechanisms that make it possible to solve complex tasks with limited resources. Another meaning of the term refers to *cognitive structures* that underlie such methods. Over the past 20 years, this has become the most common usage within the AI community, and it deserves closer inspection, for even this narrow connotation has accumulated misconceptions. I should start by stating a third claim about heuristics, in this sense of the word, that few readers will find controversial:

- *Heuristics are a type of knowledge that aids decision making and problem solving*.

As noted earlier, such knowledge can take different forms. For instance, a variety of problem-solving systems, including SAGE (Langley, 1985), Soar (Laird et al., 1987), and PRODIGY (Carbonell et al., 1990), have used rules to encode heuristics for guiding search. Larger-scale knowledge structures, such as macro-operators (e.g., Iba, 1989) and hierarchical task networks (Nau et al., 2001), can also encode heuristic content. Like search-control rules, these can make plan generation much more efficient than working with only primitive operators. However, it is important to distinguish between content that is definitional and knowledge that guides behavior. One example comes from ICARUS' (Choi & Langley, in press) distinction between concepts, which describe classes of situations, and skills, which refer to conceptual predicates but specify how to achieve goals. The former serve as definitional knowledge, whereas the latter act as heuristics to guide agent activity.

A common misconception is that heuristic knowledge must be domain specific. This holds for most examples that have appeared in the AI literature, but it is not required, and some of the most interesting cases are quite general. As Gabaldon, Langley, and Meadows (2014) have noted, the cognitive structures that guide dialogue are domain independent in the sense that they do not refer to domain predicates or the content being communicated. For instance, when we ask someone a question, we expect them to answer it or to state they do not know, independent of the details. Some evaluation functions used in planning and problem solving, such as the number of goals an operator would achieve, are equally generic. Marsella and Gratch (2009) have postulated that *emotions* modulate more basic cognitive processes for decision making and planning, with Langley (2017) arguing that the structures which underlie emotions specify abstract relations among goals, beliefs, and expectations without referring to domain-level content. These observations suggest that heuristic knowledge is still more central to intelligent behavior than originally assumed.

Unfortunately, the use of explicit knowledge structures has fallen into disfavor among many AI researchers. Game-playing systems typically use a numeric evaluation function to rank states and select among candidate moves, with structural information hidden in opaque features of states. Recent applications of neural networks in game playing (e.g., Clark & Storkey, 2015) have taken this idea still further, introducing features that are even less subject to interpretation by humans. Psychologists have long held that some forms of human expertise are 'implicit' and thus not easily communicated, but this often starts as transparent strategies which become automatized. However, few advocates of evaluation functions see benefits to associating them with explicit elements like rules or goals, making cognitive systems one of the few communities where research on heuristic knowledge structures remains a respectable topic.

6. Heuristics and Learning

The characterization of heuristics as knowledge raises natural concerns about their origin. Classical AI work on decision making, problem solving, and natural language processing assumed that humans entered these elements manually. This approach was often associated with the *expert systems* movement (e.g., Hayes-Roth, Waterman, & Lenat, 1983), but in fact it was far more widespread. Models of human expertise in cognitive psychology (e.g., Chi, Glaser, & Farr, 1988) adopted similar assumptions, as they focused on the representation and use of heuristic structures rather than on their acquisition from experience.

However, the increased excitement in recent years about machine learning suggests a final postulate that complements the ones presented earlier:

- *Heuristic knowledge structures can be acquired from experience at rapid, human-like rates.*

This statement has two components. The first is now widely accepted, at least for implicit content, but the second part of the claim differs from the mainstream view. Many researchers seem not only satisfied with, but even proud of, learning methods that require millions of training problems and thus exhibit much slower acquisition of expertise than observed in humans. In some cases, this has been motivated by (often false) rhetoric about the availability of large data sets; in others, it has been linked to guaranteed acquisition of optimal strategies, which typically holds only in the limit. Both views have drawn attention away from the key challenges.

As I have argued elsewhere (Langley, 2016), the reason for this discrepancy lies in modern machine learning's emphasis on statistical analysis. This approach can estimate the parameters of numeric evaluation functions and extract rules from observations, but it does not resemble the manner in which people acquire heuristic knowledge. Human learning is concerned primarily with the *creation of new cognitive structures*, such as control rules or HTN methods, from individual training cases, which in turn supports the rapid, incremental, and cumulative acquisition of knowledge. Statistical techniques are useful, yet they serve best not to generate alternative hypotheses, as often assumed, but rather to evaluate them (e.g., Carbonell et al., 1990). Fortunately, most cognitive systems researchers follow this approach to acquisition, which has close relatives in the literature on computational models of human learning (e.g., Anderson & Lebiere, 1998).

Also note that learning has long been viewed in terms of search through a space of structures, parameters, or their combination (Simon & Lea, 1974; Langley, Gennari, & Iba, 1987). Naturally, many researchers have drawn on heuristic methods to guide or constrain this search, some focused on generation of candidate structures and others on their evaluation. PRODIGY (Carbonell et al., 1990) adopted a hybrid approach that used explanations of problem-solving results to construct rules for search control and collected statistics on their benefits to determine which ones to retain. Langley (1995) has discussed a variety of schemes that aid incremental learning, including incorporation of background knowledge and helpful training regimes. Unfortunately, modern work on statistical induction has abandoned the most powerful idea from this older tradition – that interleaving the processes of performance and learning leads any heuristics that influence the former to constrain the latter in turn. In contrast, the cognitive systems community both respects and encourages research on this integrated approach and other ways to support rapid acquisition.

7. Conclusions and Recommendations

In the previous pages, I reviewed three different senses of the term *heuristic*, one related to methods for solving complex cognitive tasks and the others linked to structures used by such techniques. These distinctions let me state, and elaborate on, four hypotheses – about tractability, satisficing, knowledge, and learning – that clarifed the role such methods and structures play in both humans and AI systems. They also let me review the history of heuristics in artificial intelligence, which started as one of the field's most powerful, nearly defining, concepts but which has, in recent years, become maligned and misused by the AI mainstream, even though its original senses remain valid and central to the cognitive systems movement.

For this reason, it is important that members of our community take steps to counter the warped but widespread views on this important topic. This can be done most effectively in the context of their own research. To this end, I encourage authors who are reporting results to the cognitive systems community to include in their publications:

- Statements about the *role* of heuristic methods and structures in their approach to producing high-level cognition;
- Descriptions of the manner in which they *represent* heuristic knowledge, how this content is used, and how it is acquired;
- *Examples* of these cognitive structures and the ways in which they reduce cognitive load, guide problem-space search, or otherwise aid processing;
- Explicit *hypotheses* about the benefits of using these methods or structures in making complex tasks tractable even when computational resources are limited; and
- *Evidence* in the form of empirical tests, formal analyses, careful arguments, or compelling anecdotes that supports these claims about intelligent behavior.

Papers that incorporate such material will further clarify the value of heuristic approaches for the study and development of cognitive systems, as well as the conditions under which they offer benefits. In the process, they will help counter critiques that such methods are ad hoc or unfounded, which in turn will foster wider adoption of the original notion of heuristics championed by AI's founders, who I believe would support the vision of the cognitive systems movement.

Acknowledgements

This essay was supported by Grants N00014-15-1-2517 and N00014-17-1-2434 from the Office of Naval Research, which is not responsible for its contents. I thank Michael Barley, David Nicholas, and Stephanie Sage for discussions that helped refine the ideas in this paper.

References

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

Carbonell, J. G., Knoblock, C. A., & Minton, S. (1990). PRODIGY: An integrated architecture for planning and learning. In K. VanLehn (Ed.), *Architectures for intelligence*. Lawrence Erlbaum.

Chi, M. T. H., Glaser, R., & Farr, M. J. (1988). *The nature of expertise*. New York: Psychology Press.

Choi, D., & Langley, P. (in press). Evolution of the ICARUS architecture. *Cognitive Systems Research*.

Clark, C., & Storkey, A. (2015). Training deep convolutional neural networks to play Go. *Proceedings of the Thirty-Second International Conference on International Conference on Machine Learning* (pp. 1766–1774). Lille, France.

Cormen, T. H., Leiserson, C. E., & Rivest, R. L. (1990). *Introduction to algorithms*. Cambridge, MA: MIT Press.

de Groot, A. D. (1978). *Thought and choice in chess* (2nd Ed.). The Hague: Mouton Publishers.

Gentner, D., & Forbus, K. (1991). MAC/FAC: A model of similarity-based retrieval. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 504–509). Chicago: Lawrence Erlbaum.

Gabaldon, A., Langley, P., & Meadows, B. (2014). Integrating meta-level and domain-level knowledge for task-oriented dialogue. *Advances in Cognitive Systems*, *3*, 201–219.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky. *Psychological Review*, *103*, 592–596.

Gigerenzer, G. & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.

Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Hart, P. E., Nilsson, N. J., & Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics SSC4*, *4*, 100–107.

Hayes-Roth, F., Waterman, D. A. & Lenat, D. B. (Eds.) (1983). *Building expert systems*. Boston: Addison-Wesley.

Iba, G. A. (1989). A heuristic approach to the discovery of macro-operators. *Machine Learning*, *3*, 285–317.

Kaelbling, L. P. Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, *101*, 99–134.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–292.

Kirsch, A. (2011). Humanlike problem solving in the context of the traveling salesperson problem. *Advances in Cognitive Systems: Papers from the 2011 AAAI Fall Symposium* (pp. 146–153). Arlington, VA: AAAI Press.

Kirsch, A. (2017). Lessons from human problem solving for cognitive systems research. *Advances in Cognitive Systems*, *5*, 13–24.

Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, *33*, 1–64.

Langley, P. (1985). Learning to search: From weak methods to domain-specific heuristics. *Cognitive Science*, *9*, 217–260.

Langley, P. (1995). Order effects in incremental learning. In P. Reimann & H. Spada (Eds.). *Learning in humans and machines: Towards and interdisciplinary learning science*. Oxford: Elsevier.

Langley, P. (2016). The central role of cognition in learning. *Advances in Cognitive Systems*, *4*, 3–12. .

Langley, P. (2017). A cognitive systems analysis of personality and conversational style. *Proceedings of the Fifth Annual Conference on Cognitive Systems*. Troy, NY.

Langley, P., Gennari, J. H., & Iba, W. (1987). Hill-climbing theories of learning. *Proceedings of the Fourth International Workshop on Machine Learning* (pp. 312–323). Irvine, CA: Morgan Kaufmann.

Levesque, H. J., & Brachmann, R. J. (1987). Expressiveness and tractability in knowledge representation and reasoning. *Computational Intelligence*, *3*, 78–93.

Marsella, S. C., & Gratch, J. (2009). EMA: A process model of appraisal dynamics. *Cognitive Systems Research*, *10*, 70–90.

Nau, D. S., Cao, Y., Lotem, A., & Muñoz-Avila, A. (2001). The SHOP planning system. *AI Magazine*, *22*, 91–94.

Newell, A. (1982). The knowledge level. *Artificial Intelligence*, *18*, 87–127.

Newell, A. (1983). The heuristic of George Polya and its relation to artificial intelligence. In R. Groner, M. Groner, & W. F. Bischof (Eds.), *Methods of heuristics*. Hillsdale, NJ: Lawrence Erlbaum.

Newell, A., Shaw, J. C., & Simon, H. A. (1957). Empirical explorations of the Logic Theory Machine. A case study in heuristic. *Proceedings of the Western Joint Computer Conference* (pp. 218–230) New York: Institute of Radio Engineers.

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

Polya, G. (1957). *How to solve it*. Garden City, NY: Doubleday.

Selz, O. (1927). Die umgestaltung der grundanschauungen vom intellektuellen geschehen [The revision of the fundamental conceptions of intellectual processes]. *Kantstudien*, *32*, 273–280.

Simon, H. A. (1947). *Administrative behavior: A study of decision-making processes in administrative organization*). New York: Macmillan.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*, 129–138.

Simon, H. A. (1981). Otto Selz and information-processing psychology. In N. H. Frijda & A. D. de Groot (Eds.), *Otto Selz: His contribution to psychology*. The Hague: Mouton.

Simon, H. A., & Lea, G. (1974). Problem solving and rule induction: A unified view. In L. W. Gregg (Ed.), *Knowledge and cognition*. Potomac, MD: Lawrence Erlbaum.

Slagle, J. R. (1971). *Artificial intelligence: The heuristic programming approach*. New York: McGraw-Hill.

ter Hark, M. (2010). The psychology of thinking before the cognitive revolution: Otto Selz on problems, schemas, and creativity. *History of Psychology*, *13*, 2–24.

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

Weber, E. U., & Johnson, E. J. (2009). Mindful judgment and decision making. *Annual Review of Psychology*, *60*, 53–85.