Advances in Cognitive Systems 7 (2018) 13–22 Submitted 12/2018; published 12/2018

Planning Systems and Human Problem Solving

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

Early AI planning systems borrowed key ideas from studies of human problem solving, but the past 30 years have seen researchers largely abandon this strategy. This has led to computational methods that are highly efficient for certain classes of problems, but that lack many desirable features of people's cognition. In this paper, I review the initial interactions between the two fields, the reasons for their gradual separation, and the potential benefits of renewed interchange. I examine abilities that human problem solvers exhibit but that receive scant attention in the planning literature, along with ideas for addressing these omissions. In closing, I suggest ways to encourage more research in AI's original interdisciplinary tradition, which the cognitive systems paradigm adopts.

1. Introduction

The ability to generate extended plans that achieve goals is a distinguishing feature of human intelligence. Other organisms carry out actions that produce desirable states, but they have very limited foresight, whereas people excel at long-range reasoning about the means to achieve their ends. This makes planning and problem solving natural areas of study for artificial intelligence and cognitive psychology, and there has been substantial work in each tradition on similar tasks and issues. Since the 1950s, research in these complementary disciplines has produced many insights about the structures and processes that underlie this important cognitive capacity.

In this paper, I examine the multifaceted relationship between these interconnected scientific movements. First I review historical links between the planning community and cognitive psychology, including some of the reasons they grew apart. After this, I examine seven important characteristics of human problem solvers that could drive AI planning efforts in fruitful directions. The purpose is not to replicate the details of problem solving in our species, but rather to use it as a source of capabilities and mechanisms that could improve planning systems. Finally, I discuss ways to foster research that addresses these important but overlooked topics, along with some benefits of renewed interchange between the two intellectual communities.

2. Historical Background

Early research on AI planning was strongly influenced by findings on human problem solving. Heuristic search first appeared in the Logic Theorist (Newell et al., 1958), which modeled how subjects proved theorems in propositional logic, and the earliest planning system, GPS (Newell et al., 1960), incorporated insights from think-aloud protocols. Developments in knowledge-based systems occurred in tandem with studies of human expertise in domains like chess (de Groot, 1976), while results about human improvement on problem-solving tasks (Anzai & Simon, 1979) influenced early work on machine learning. Both the AI community and cognitive psychology gained substantially from this interaction, with the former extracting ideas for automated planning and the latter obtaining promising models for high-level human behavior.

Despite these mutual benefits, the situation has changed drastically over the past 30 years. During this period, the community has shifted its focus to planning methods that offer formal guarantees and solve certain types of tasks efficiently. Although researchers have retained their original concern with domain-independent approaches, they have also introduced strong assumptions that limit the generality of their results. The same period has seen decreased interest in findings about human problem solving and increased doubts about their usefulness. In many ways, the computational study of planning has become the polar opposite of the vision originally proposed by AI's founders.

Both technical and sociological reasons lie behind this separation. Increased emphasis on formal analysis, due to AI researchers' training in computer science, has produced concerns about completeness, correctness, and optimality that lead many to view psychological findings as irrelevant. Introduction of quantitative performance metrics and competitions has encouraged incremental progress on standardized problems, rather than efforts to reproduce the full range of human abilities. Increased computer speed and storage capacities have enabled simple-minded CPU-intensive and memory-based approaches, such as grounding operators at compile time, that diverge from human processing. Finally, successful applications on narrowly defined problem classes have fostered research on similarly limited tasks, rather than on the general capabilities observed in people.

Taken together, these factors have led many AI researchers, including those focused on planning, to eschew results from psychology and to look elsewhere for inspiration. Nevertheless, there are good reasons to resurrect the early alliance between the two disciplines. Despite advances in planning methods that have produced superhuman results on some domains, there are ways they fall short of human intelligence that are worth addressing in future research. Again, the aim is not to reproduce the details of problem solving in people, but to use insights from it as means to the end of better planning systems. Human behavior is not the only source of such ideas, but it is an important and valuable one that the research community would be foolish to ignore.

3. Characteristics of Human Problem Solving

In this section, I consider some key aspects of human cognition that are not emphasized by the modern AI planning community. In each case, I review the main psychological findings, discuss related research, clarify how current planning efforts fail to make contact with these ideas, and suggest ways to remedy the situation. Both Fahlman (2012) and Kirsch (2017) have made similar points about useful lessons that studies of human problem solving offer to cognitive systems researchers.

3.1 Satisficing and Heuristic Methods

Both everyday experience and empirical studies show that human decision makers do not always make 'optimal' selections. Simon (1956) argued that they *satisfice*, in that they find alternatives which are *good enough*. When renting an apartment or buying a refrigerator, we do not examine all choices in detail; instead, we consider subsets that we find manageable and halt on finding acceptable ones. Simon connected satisficing to the notion of an *aspiration level*, which informs us when an option is sufficient. Other researchers (e.g., Gigerenzer & Todd, 1999; Weber & Johnson, 2009) have collected substantial evidence for satisficing behavior in a wide range of settings.

Heuristic processing was a defining feature of early AI research (Langley, 2017) that distinguished it from the field of computer science, which focused on algorithmic methods. Satisficing was a recurring theme in work on intelligent systems until the 1980s, when computer science departments, with their emphasis on formal analyses, became the default home for AI scientists. This encouraged planning researchers to insist on features like soundness, completeness, and optimality, especially after introduction of A* (Hart et al., 1968) showed the latter was possible with heuristic methods. Although work on satisficing continues and even has a place in the International Planning Competition, there remains a strong bias toward planners that find optimal solutions and, for decades, convergence to optimal control policies was a defining feature of reinforcement learning.

The drawback is that, by restricting research to methods that are optimal, or even complete, the community fails to consider approaches that do not have these features but that generally work well on a wider class of tasks, as do human problem solvers.[1] There is nothing wrong with using A* or even exhaustive search for some problems, but they are clearly intractable for other tasks, and we should not restrict our efforts to the these schemes. Moreover, the definition of optimality becomes unclear when an agent must determine its own objectives (Aha et al., 2013), yet we can still develop systems that combine goal reasoning with plan generation. The planning community would benefit from returning to AI's original views on heuristic methods and recognizing the value of work in this paradigm. We should encourage more researchers to explore mechanisms that, like people, are effective in practice despite lacking formal guarantees.

3.2 Hierarchical Problem Decomposition

Verbal protocols suggest that even novices often solve problems by decomposing complex tasks into simpler ones and, if they are successful, combining the results into overall solutions (Newell & Simon, 1972). There may be many possible decompositions, only some of which are viable, but, when combined with heuristics, search through this space can be very effective. Early studies of problem solving revealed that, for some tasks, humans often use means-ends analysis, which selects an operator to achieve some goal and decomposes the current problem on that basis. GPS (Newell et al., 1960) reproduced this aspect of human cognition first, and STRIPS (Fikes et al., 1972) extended the method to operate over more transparent representations. The PRODIGY architecture (Carbonell et al., 1990) also incorporated means-ends analysis as a central mechanism.

However, because the standard version of this method has difficulty on tasks that involve goal interactions, the AI community had largely abandoned the approach by the 1980s, focusing first on partial-order methods and later on forward-chaining techniques, which currently dominate the literature. A few researchers have extended the earlier technique to overcome its limitations (Borrajo & Veloso, 1994; Jones, 1993), while others (Marsella & Schmidt, 1993; Ruby & Kibler, 1993) have developed different mechanisms for searching a space of problem decompositions. However, very few results in the wider planning community have built on their promising contributions.

1. Research on cognitive architectures (e.g., Laird, 2012; Choi & Langley, 2018) continues in the original satisficing tradition, but work in this area has had reasonably little impact on the rest of the field.

Regardless, humans rely on means-ends analyis and similar methods for an obvious reason. Breaking problems into subproblems, when successful, decreases the effective depth of solutions and thus reduces exponentially the overall amount of search. The result is similar to abstraction planning, another idea with roots in studies of human cognition (Newell & Simon, 1972), which can transform otherwise intractable tasks into manageable ones. This does not mean that solving problems by decomposing them is the only viable response to complexity, but it does offer a promising approach to difficult tasks. Thus, problem decomposition merits more attention from modern AI researchers, who should broaden their own search efforts through the space of planning strategies.

3.3 Knowledge Lean and Rich Processing

Early studies of human problem solving addressed novice behavior on reasonably simple tasks, such as puzzles, which led to key insights about problem-space search (Newell et al., 1960). Later work examined expert behavior on more challenging tasks, revealing the importance of knowledge in limiting or even eliminating search (Larkin et al., 1980). However, psychological research on chess (de Groot, 1978) and other complex domains clarified that human problem solvers often combine these abilities, using knowledge when available but resorting to search when needed.

The AI planning community has made quite substantial progress on some of these elements but not others. Most work has emphasized search in knowledge-lean settings, relying on heuristic search to generate plans from primitive operators. A smaller group has focused on hierarchical task networks (Nau et al., 2003), which encode domain knowledge as a set of methods that selectively decompose tasks into simpler subtasks. HTN planners combine knowledge with search, but few systems within this paradigm can fall back on primitive operators when no such knowledge is applicable, which is a definite feature of human problem solving. The advantage of this latter capability should be clear. Knowledge enables far more effective planning, but it can be incomplete, in which case resorting to a search-intensive method is a natural response.

Research on problem-solving architectures, such as Soar (Laird, 2012) and PRODIGY (Carbonell et al., 1990), fare much better on this dimension, as they use control rules to guide search when applicable but invoke fuller exploration when not. But this work has had little influence on the broader AI planning community, and only recently has the HTN paradigm started to combine the two strategies (e.g., Shivashankar et al., 2013; To et al., 2015). Such research remains atypical, which suggests that we need far more effort to integrate these two important facets of problem solving. This would let planning systems take advantage of knowledge to solve tasks when it is available but avoid fragility when these structures are absent or incomplete.

3.4 Rapid Acquisition of Cognitive Structures

Human problem solvers not only learn from their experiences; they acquire new cognitive structures very rapidly. People gain knowledge about a domain from their own search efforts (Anzai & Simon, 1979), from sample solutions (VanLehn & Jones, 1993), and, of course, from direct instruction. They also appear to collect statistics, more gradually, about the usefulness of these structures (Anderson, 1993), but they create the knowledge elements themselves from experience with individual tasks. These incremental processes underlie humans' ability to master new skills in mathematics, physics, chess, and other complex domains with impressive effectiveness.

This contrasts sharply with most work on learning for planning and control. For example, many techniques for reinforcement learning operate incrementally, but they typically require hundreds of times more training problems than humans. Systems that use such techniques may ultimately acquire policies that outperform people, but this can take millions of trials (Silver et al., 2016). Despite its reliance on statistical inference and inefficient use of data, this paradigm and its relatives remain very popular in the planning community. Nonincremental methods for inducing control knowledge for action selection are more rapid, but they still lack the speed of humans and the ability to learn new structures from experience with individual training problems.

Some earlier research on planning and learning comes much closer to exhibiting human characteristics. For example, Soar (Laird et al., 1987) and PRODIGY (Carbonell et al., 1990) acquired control rules by analyzing successes and failures during search on individual problems. More recently, Nejati et al. (2006) and Hogg et al. (2008) reported methods for learning HTNs analytically from sample solutions, while Li et al. (2012) used a similar approach for learning from successful search. However, such results are the exception rather than the rule in AI planning circles, which favor statistical methods that master domains far more slowly than humans, and the field would benefit from a more balanced portfolio. Even methods for reinforcement learning could benefit from simple forms of analysis. Given a factored reward function and models of actions' effects, they could reason about how each action sequence influenced some elements but not others, giving more effective credit assignment and increasing the rate of learning.

3.5 Strategic Variation and Adaptation

Despite high-level regularities in human problem solving, such as reliance on heuristic search and use of knowledge to reduce it, people also show considerable variety in their strategic behavior. For instance, they typically use means-ends analysis on puzzles like the Tower of Hanoi (Anderson, 1993), but rely on forward search in games like chess (de Groot, 1978). They often halt when they find a single solution, but in some cases obtain multiple answers and select the best one. On simple games like Tic-Tac-Toe, they may carry out exhaustive depth-first search, whereas on complex ones like chess, they use methods like progressive deepening with fewer memory demands. More important, they make these strategic decisions on their own, sometimes on an individual problem. For instance, Simon and Reed (1976) observed strategy adaptations on subjects given the Missionaries and Cannibals task. This variability and adaptability is seldom discussed in the psychological literature, but its existence is widely accepted, and it presumably occurs for some reason.

The AI planning community has devoted little attention to these abilities. Many systems include parameters that control details of plan generation, but usually within a single strategic framework, such as forward search. Soar (Laird et al., 1987) could mimic many different search methods, while FLECS (Veloso & Stone, 1995) supported both backward and forward chaining, but neither emphasized strategic adaptation. Portfolio methods for planning invoke a number of search techiques in parallel and return the first solution found, but this is not human-like adaptation. Techniques for automated configuration of planners select a combination of parameters within a framework, but they tune on a set of training problems, which comes closer to learning. Felner et al. (2010) describe a bidirectional search method that chains in the direction with lower branching factor, while Langley et al. (in press) report a more flexible approach that adapts the direction of chaining, amount of backtracking, and termination criterion, but these efforts remain only a start.

In summary, the AI literature contains few examples of flexible problem solvers, and truly adaptive ones are even less common. Humans appear to adopt search strategies relevant to the problem at hand not by selecting from a fixed set of methods, but rather by composing them from more basic elements. If we want AI planning systems to incorporate the same abilities, researchers should shift their efforts away from a few narrow paradigms and devote more energy to understanding the entire space of planning methods. The community should devise theories of problem solving that support the entire range of known search methods and attempt to discover when and why some strategies are more effective than others. Incorporating these findings as heuristics to guide adaptation may let us create planners that are as versatile and adaptable as people.

3.6 Creativity and Problem Reformulation

Human problem solving often involves routine search, but in some cases people exhibit creativity by finding solutions that are far from obvious. There have been many psychological studies of such insights and the factors that influence them, but only a few theoretical accounts have been stated in computational terms. Perhaps the most plausible explanation is that insight involves a shift in problem formulation. A classic example is the mutilated checkerboard (Kaplan & Simon, 1990), where a representation change makes a difficult problem almost trivial, and Ohlsson (1992) offers a similar analysis of other insight tasks. Both treatments posit that the key process involves some change to how one encodes states, goals, or operators.

There has been little research within the AI community on this topic, possibly due to a widespread belief that, because automated planners can search more extensively than humans, it is unnecessary. However, Riddle et al. (2016) have shown that changes to problem formulations can transform tasks that are intractable for even the most advanced planner into ones that are easily solvable. Building on Amarel's (1968) early analysis of representation change in problem solving, they report a system that applies meta-level operators to automatically shift its representation of planning tasks, which often leads to major reductions in the search carried out.

These recent results are encouraging, but creativity in problem solving deserves far more resources than it has been allocated. Computational methods for problem reformulation and representation change would produce systems that are less dependent on careful human crafting, making them as robust as the best human planners in their ability to solve difficult problems. Sarathy and Scheutz (2018) provide a formal analysis of such creative problem solving and propose a promising research agenda for addressing it more fully. We need more research on computational artifacts that exhibit this ability, whether or not they achieve creativity in the same manner as humans.

3.7 Explanation and Justification

Human planners can often explain why they consider different alternatives and how they decide which ones to pursue. Early studies of problem solving relied heavily on verbal protocols (Newell & Simon, 1972), in which subjects say their thoughts aloud for later analysis. Retrospective reports are less reliable than ones generated during the search process, but the ability to produce them at all shows that people have at least partial episodic traces of decision making. They can retrieve and use this stored information to explain the reasoning behind their plans and activities. Moreover, they can link their decisions not only to personal goals and evaluation criteria; they can also justify them in terms of social norms that reflect laws, customs, and cultural mores (Malle et al., 2015)

The outputs of most AI planning systems already record the causal links among states, actions, and goals, but few of them retain this content across successive tasks. A few analogical planners, like PRODIGY (Borrajo & Veloso, 1994) and EUREKA (Jones, 1993), store, index, and retrieve details about which branches of the search tree produced success or failure. However, they use this information when solving future problems, as do other systems that incorporate episodic memories (e.g., Laird, 2012). They do not access it to support retrospective explanation of their decision making and, although work on 'explainable AI' is currently in vogue, only a few researchers have addressed the challenge of *explainable planning* (Fox et al., 2017; Smith, 2012).

Thus, replicating the ability to store, index, and retrieve problem-solving traces for use in explanation remains an important challenge for the community. We need more research on this important topic and, as in other areas, studies of human cognition offer an excellent source of ideas for how to proceed, although not the only one. Our aim should not be to reproduce the details of retrospective memory in people, which can be incomplete and error prone, but rather to combine the structures and processes that appear responsible for it with the more complete and reliable memory stores of AI systems. Moreover, we should find ways to represent and use societal norms during plan generation. Finally, we should combine the two abilities to support explanations that incorporate these norms, producing what I have elsewhere called *justified agency* (Langley, 2019).

3.8 Summary Remarks

In summary, despite important and impressive advances in technologies for automated planning, there remain many features of human problem solving that the field has placed on the sidelines. Indeed, there is less work on these topics than in AI's earliest days, when it was common for researchers to incorporate results from psychology. This does not mean that the planning community should abandon the insights and techniques developed over the past 30 years, but it does suggest its spotlight has become too narrow and that the discipline would benefit from widening the beam.

Of course, this discussion does not exhaust the characteristics of human problem solving. For instance, protocols indicate that people have finer-grained control over search than most AI planners; they apply one operator at a time, rather than expanding a node to generate all successors. This means that attention, another mechanism studied by psychologists, plays an important role in guiding the search for solutions. Emotions also influence decision making and, rather than being the undesirable factor often assumed, can direct attention in useful ways (Simon, 1967). Planning researchers should explore these and other features of human processing.

4. Bridging the Conceptual Divide

There would be multiple benefits to rapprochement between the AI planning community and studies of human problem solving. As we have just seen, the latter is an important source of challenges for the former which could drive research in interesting directions that might otherwise remain unexplored. Moreover, findings about human cognition can suggest constraints on system design, including how to represent and organize knowledge, how to use that knowledge in planning, and how to acquire knowledge from experience. Homo sapiens remain our only examples of general intelligent systems, and insights about their operation deserve serious consideration. Of course, renewed interchange could also help us understand the nature of human intelligence, which comprises an important set of phenomena that demand scientific explanation.

If we assume that bridging the conceptual divide is desirable, then we should also consider how to achieve this objective. One response is to provide broader training in AI that reviews concepts from cognitive psychology, their role in early planning research, and the challenges that they still offer. We should also encourage and publish more exploratory work that aims to reproduce abilities observed in humans without requiring that they compete with mature systems designed with other criteria in mind. In addition, agencies can foster planning research that incorporates insights from psychology; the US Office of Naval Research has long supported work in this vein, but additional funding would encourage more work that crosses the disciplines' boundaries.

In the preceding pages, I recounted how early AI planning research benefited from psychological studies but later shifted away from this tradition. This led the community to develop efficient techniques for certain classes of problems, but only at the price of ignoring important aspects of human cognition. These included satisficing and heuristic methods, hierarchical problem decomposition, combining knowledge with search, rapid learning of cognitive structures, strategic variation and adaptation, creativity through reformulation, and retrospective explanation. I also noted ways the community would gain from reestablishing its connections to cognitive psychology, along with steps it should take to achieve this aim. Without such effort, AI planning research seems likely to remain a narrow discipline that bears little resemblance to its original vision.

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 Will Bridewell, Ben Meadows, and Stephanie Sage for discussions that helped refine the ideas in this paper.

References

Aha, D. A., Cox, M. T., & Muñoz-Avila, H. (2013). (Eds.). *Goal reasoning: Papers from the ACS workshop*. Baltimore, MD.

Amarel, S. (1968). On representations of problems of reasoning about actions. In D. Michie (Ed.), *Machine Intelligence 3*, 131–171. New York: American Elsevier.

Anderson, J. R. (1993). Problem solving and learning. *American Psychologist*, *48*, 35–44.

Anzai, Y., & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review*, *86*, 124–140.

Borrajo, D., & Veloso, M. M. (1994). Incremental learning of control knowledge for nonlinear problem solving. *Proceedings of the Seventh European Conference on Machine Learning* (pp. 64–82). Catania, Italy.

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

Choi, D., & Langley, P. (2018). Evolution of the ICARUS architecture. *Cognitive Systems Research*, *48*, 25–38.

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

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

Felner, A., Moldenhauer, C., Sturtevant, N. R., & Schaeffer, J. (2010). Single-frontier bidirectional search. *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence* (pp. 59–64). Atlanta, GA: AAAI Press.

Fikes, R. E., & Nilsson, N. J. (1972). STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, *2*, 189–208.

Fox, M., Long, D., & Magazzeni, D. (2017). Explainable planning. *Proceedings of the IJCAI-17 Workshop on Explainable AI* (pp. 24–30). Melbourne.

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*, *SSC-4*, 100–107.

Hogg, C., Muñoz-Avila, H., & Aha, D. (2008). HTN-MAKER: Learning HTNs with minimal additional knowledge engineering required. *Proceedings of the Twenty-Third National Conference on Artificial Intelligence* (pp. 950–956). Chicago: AAAI Press.

Jones, R. M. (1993). Problem solving via analogical retrieval and analogical search control. In A. L. Meyrowitz and S. Chipman (Eds.), *Foundations of knowledge acquisition: Machine learning*. Boston: Kluwer Academic.

Kaplan, C. A., & Simon, H. A. (1990). In search of insight. *Cognitive Psychology*, *22*, 374–419.

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

Laird, J. E. (2012). *The Soar cognitive architecture*. Cambridge, MA: MIT Press.

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

Langley, P. (2017). Heuristics and cognitive systems. *Advances in Cognitive Systems*, *5*, 3–12.

Langley, P. (2019). Explainable, normative, and justified agency. *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*. Honolulu, HI: AAAI Press.

Langley, P., Barley, M., & Meadows, B. (in press). Adaptive search in a hierarchical problem-solving architecture. *Advances in Cognitive Systems*.

Larkin, J. H., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Expert and novice performance in solving physics problems. *Science*, *208*, 1335–1342.

Li, N., Stracuzzi, D. J., & Langley, P. (2012). Improving acquisition of teleoreactive logic programs through representation extension. *Advances in Cognitive Systems*, *1*, 109–126.

Malle, B. F., Scheutz, M., & Austerweil, J. L. (2015). Networks of social and moral norms in human and robot agents. *Proceedings of the International Conference on Robot Ethics* (pp. 3–17). Lisbon, Portugal.

Marsella, S. C., & Schmidt, C. F. (1993). A method for biasing the learning of nonterminal reduction rules. In S. Minton (Ed.), *Machine learning methods for planning* (pp. 499–535). San Francisco, CA: Morgan Kaufmann.

Nau, D., Au, T., Hghami, O., Kuter, U., Murdock, J., Wu, D., & Yaman, F. (2003). SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research*, *20*, 379–404.

Nejati, N., Langley, P., & Könik, T. (2006). Learning hierarchical task networks by observation. *Proceedings of the Twenty-Third International Conference on Machine Learning* (pp. 665–672). Pittsburgh.

Newell, A., Shaw, J. C., & Simon, H. A. (1958). Elements of a theory of human problem solving. *Psychological Review*, *65*, 151–166.

Newell, A., Shaw, J. C., & Simon, H. A. (1960). Report on a general problem-solving program for a computer. *Proceedings of the International Conference on Information Processing* (pp. 256–264). UNESCO House, France.

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

Ohlsson, S. (1992). Information processing explanations of insight and related phenomena. In M. Keane and K. Gilhooly (Eds.), *Advances in the psychology of thinking* (Vol. 1). London: Harvester-Wheatsheaf.

Riddle, P., Douglas, J., Barley, M., & Franco, S. (2016). Improving performance by reformulating PDDL into a bagged representation. *Proceedings of the Eighth Workshop on Heuristics and Search for Domain-Independent Planning* (pp. 28–36). London.

Ruby, D., & Kibler, D. (1993). Learning recurring subplans. In S. Minton (Ed.), *Machine learning methods for planning*. San Francisco, CA: Morgan Kaufmann.

Sarathy, V., & Scheutz, M. (2018). MacGyver problems: AI challenges for testing resourcefulness and creativity. *Advances in Cognitive Systems*, *6*, 31–44.

Shivashankar, V., Alford, R., Kuter, U., & Nau, D. (2013). The GoDeL planning system: A more perfect union of domain-independent planning and hierarchical planning. *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence* (pp. 2380–2386). Beijing.

Silver, D. et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, *529*, 484–489.

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

Simon, H. A. (1967). Motivational and emotional controls of cognition. *Psychological Review*, *74*, 29–39.

Simon, H. A., & Reed, S. (1976). Modeling strategy shifts in a problem-solving task. *Cognitive Psychology*, *8*, 86–97.

Smith, D. E. (2012). Planning as an iterative process. *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence* (pp. 2180–2185). Toronto: AAAI Press.

To, S. T., Langley, P., & Choi, D. (2015). A unified framework for knowledge-lean and knowledge-rich planning. *Proceedings of the Third Annual Conference on Cognitive Systems* (pp. 1–17). Atlanta, GA: Cognitive Systems Foundation.

VanLehn, K., & Jones, R. M. (1993). Integration of analogical search control and explanation-based learning of correctness. In S. Minton (Ed.), *Machine learning methods for planning*, 273–315. San Mateo, CA: Morgan Kaufman.

Veloso, M. M., & Stone, P. (1995). FLECS: Planning with a flexible commitment strategy. *Journal of Artificial Intelligence Research*, *3*, 25–52.

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