Problem Solving: Phenomena in Search of a Thesis

Bruce D. Burns (burnsbr@msu.edu)
Department of Psychology, Michigan State University
East Lansing, MI 48824-1117, USA

Regina Vollmeyer (vollmeye@rz.uni-potsdam.de)
Institut für Psychologie, Universität Potsdam, Postfach 601553
14415 Potsdam, Germany

Abstract

Problem solving has been the study of a set of phenomena rather than a set of theories. Newell & Simon's (1972) concept of search has proved very useful for describing problem solving but it is not a testable theory. We point out that without testable theories, thought about problem solving cannot progress through the interaction of thesis and antithesis. Problems solving requires theories and we propose a specific form of multispace search theory. The hierarchical three-space theory of problem solving can be derived from existing literature, and proposes that interactive search occurs in instance space (states of the problem), rule space (possible rules that govern the problem), and model space (the general understanding of the problem). This theory could be used to generate testable predictions regarding the interaction of spaces and provides a way to try to unify diverse phenomena.

Problem Solving: What is the theory?

Sixty years ago, Maier (1940) noted that problem solving was frequently cited as a barren field of inquiry. He argued that this blandness is due to the proliferation of experimental tasks which render generalization difficult. It is bland because there is a set of phenomena, but no underlying explanation of them.

Has this blandness diminished since Maier wrote this? If one examines the way the field presents itself to its first line of consumers, undergraduates taking introductory courses in cognitive psychology, then it is arguable that the same problems identified by Maier (1940) continue to bedevil problem solving research. When one picks up a typical introductory cognitive psychology book and turns to the sections on perception, attention, or memory, then one find a lively description of competing theories and the evidence used to support/discredit them. These are on-going debates so different books present these debates in different ways depending on the biases of the author. In contrast, there is a remarkable similarity between different books when one turns to the section on problem solving. Anderson (2000) is a typical example. He covers procedural knowledge and search, operators (including analogy), operator selection (including means-ends analysis and the Tower of Hanoi), representation (including functional fixedness), and set effects. Other books may differ in their details, but cover the same basic ground. What is noticeable by their absence, are theories of problem solving.

There has been progress since 1940, in particular Newell and Simon's (1972) idea of problem solving as search of a problem-space. This has been very valuable both for psychological and computational approaches to problem solving, but it is not a complete theory (as others have also noted, such as VanLehn, 1989). As a language for thinking about problem solving, search has proved to be useful and enduring. However, it makes no testable predictions, so there can be no competing theory.

Sternberg (1995) emphasizes the dialectical progression of ideas in psychology. As described by Hegel (1807/1931), the dialectic begins with a viewpoint that is proposed and believed (a thesis), but in response a competing view arises (an antithesis), and eventually the best features of both are melded into a synthesis. Then the process starts again. Sternberg argues that much of the history of psychology can be seen in terms of the dialectic. This progression cannot occur in the field of problem solving because there is no thesis, therefore there can be no antithesis, and there can be no synthesis. Anderson (2000) and other introductory cognitive psychology books illustrate that we know quite a lot about the phenomena of problem solving, but there is no thesis for the phenomena. VanLehn (1989) provides a list of robust empirical findings regarding problem solving. (Space precludes us from taking on a related difficulty with problem solving research, its definition. The definition has varied from very narrow [essentially the study of solving the Tower of Hanoi] to the very broad [every goal oriented activity]. For the purpose of this paper we use VanLehn's characterization of problem solving as multistep goal-directed tasks that last a few minutes to a few hours.)

The aim of this paper is to try to present a thesis, or to at least struggle towards that aim. In doing this we do not wish to throw away the power of treating problem solving as search, instead we want to try to wield it into a form that presents a testable thesis. We do this by taking seriously another part of Maier's (1940) paper, that problem solving is not a single process, but a set of processes. This leads to multispace search theories, and we propose a specific form of a hierarchical three-space search theory of problem solving. As may already be apparent this paper is speculative and, merely uses existing ideas, but we wish to show how these ideas can be put together in such a way that a hierarchical three-space theory falls out.

Multispace Search

Problem Solving as Search. Newell and Simon (1972) proposed that for every problem there exists a problem space which is defined by three components: 1) the *initial state* of the problem; 2) a set of *operators* that can transform a problem state; 3) a *test* for whether a problem state constitutes a solution (this may be a particular *goal state* or set of goal states). Finding a solution is a process of searching the set of states logically defined by the initial state and the operators that can be applied, until a solution is found. This terminology has proved to be useful for describing a wide range of problem solving behavior. However, to encompass a wider range of phenomena, this framework has been extended in two ways. In order to include induction and problem solving within the same framework, Simon and Lea (1974) claimed that search occurs in a *dual-space*. In order to capture the influence of different representations, Hayes and Simon (1974) claimed that an *understanding process* is required as well as a search process.

Dual-Space Search. Simon and Lea (1974) proposed that problem solving does not necessarily consist of search of a *single* problem space. To encompass multiple spaces they generalize the components of Newell and Simon's (1972) description of problem solving in the following way: 1) the elements of a problem space are *knowledge states*; 2) operators are *generative processes* that take a knowledge state as input and produce a new knowledge state as output; 3) there are one or more *test processes* for determining solution and for comparing knowledge states; 4) there are *selection processes* for which of these generators and tests to employ, on the basis of the information contained in the knowledge states. Induction can then be related to problem solving by allowing a dual-space search to be conducted. The search for rules that describe a task is conducted in a *rule space*, the states of which are all possible rules, and the operators are processes for generating, modifying and testing rules. Testing, however, requires movement within *instance space*, consisting of all possible states of the task, and the operators are processes allowed by the task for moving between instance states. Thus the two problem spaces are conceptually distinct, but intimately related; test processes for rule space lead to the generation of instances, whereas information that results from such instances leads to movement in rule space.

Simon and Lea (1974) suggested that many induction tasks can be described in this dual-space framework. For example, in concept attainment tasks learners generate possible rules from instances. They then test or select between alternative rules by observing or creating relevant instances. Thus concept attainment can be seen as a dual-space search, in which the goal is in rule space.

Problem solving may involve search of instance space only, but it is a dual-space search if a problem solver tries to learn rules which can be generally applied to reaching different goal states. Simon and Lea (1974) pointed out that the Tower of Hanoi problem is usually thought of as search of instance space: find a sequence of moves that transfers all disks to the goal peg. This is a single space search. But the task could be described as: find a *rule* for transferring disks from one peg to another (e.g., the first move depends on whether the number of disks is odd or even). This requires a dual-space search.

Dual-space search has been extended to scientific reasoning by Klahr and Dunbar (1988) in their Scientific Discovery as Dual Search (SDDS) model. They proposed that in scientific reasoning people have an *hypothesis space* (similar to rule space) and an *experiment space* (similar to instance space). Reasoners propose hypotheses, and test them by conducting experiments. Klahr and Dunbar (1988) and Dunbar (1993) found that subjects who tested hypotheses performed a learning task better, a finding which supports SDDS.

More evidence that dual-space search occurs in problem solving was found by Vollmeyer, Burns, and Holyoak (1996). They had participants learn to control a complex system called *biology-lab* in which they could manipulate inputs and observe the changes to the outputs. Ultimately they had to bring the system to a set of output states, but participants were not told the nature of the set of equations linking inputs to outputs. Vollmeyer et al. manipulated the goals of problem solvers by either telling them what the goal state was before they started exploring the system (a specific goal), or delaying informing them of the goal until after they had explored (a nonspecific goal). They found that a group given a specific goal learned less about the structure of the biology-lab task and transferred more poorly to a new goal, than did the nonspecific goal group. The strategies of the specific goal group indicated search of instance space (i.e., find a path to the goal), but the nonspecific goal group instead appeared to test rules (i.e., search hypothesis space).

Understanding Processes. Before a problem solver can attempt a problem, the problem instructions must be understood. The importance of understanding processes in natural language has been well illustrated (e.g., Bransford & Franks, 1971). Hayes and Simon (1974, 1977) explored the impact that understanding can have on problem solving. Hayes and Simon (1977) gave subjects different isomorphic versions of the Tower of Hanoi problem and found a dramatic effect on solution ease from the ease of understanding the problem description. The importance of representation in problem solving was a point emphasized long ago by the Gestaltists (e.g., Maier, 1930).

Hayes and Simon (1974) incorporated understanding into Newell and Simon's (1972) framework by proposing understanding as a subprocess that cooperates with search of the problem space. The search process is driven by the result of understanding processes, rather than the problem itself. However, it may not be that understanding processes first produce a representation of the problem, and then search takes over. The two processes may alternate or even blend together (see Hayes & Simon, 1974). That representations may be fluid and interact with attempts to solve a problem, is a point also made by researchers working within other frameworks (e.g., Burns, 1996; Hofstadter, 1995).

A Hierarchical Three-space Theory

Integrating Understanding Processes. If understanding processes create the representation of problem space, then in a dual-space search theory these processes must create the representation of not only the instance space, but of the rule space as well. Thus understanding processes define the instance states that can be searched, and do so via the candidate rules that might govern instance states. The research on functional fixedness (e.g., Maier, 1930) can be seen in terms of the problem solver's understanding processes defining the wrong rule space. Similarly, the research on how false assumptions can be a barrier to solution can be viewed in this way. Weisberg and Alba (1980) showed that problem solvers attempting the nine-dot problem could only solve it when their assumption that they could only draw lines within a restricted area was removed. In our terms, they were searching the wrong rule space. Of course, having the correct rule space does not guarantee success (as Weisberg & Alba found) as having the correct rule space to search is not equivalent to having the correct rule.

Given that representations may change during problem solving, understanding processes can be seen as conducting a form of search. VanLehn (1989) suggested that schema selection can be a form of search when a person is uncertain as to which schema to select. For example, Larkin (1983) gave expert physicists a straightforward, but difficult, physics problem to solve. Although two of the five physicists immediately selected the correct schema for solving the problem, the other three physicists tried two or more schemas. In this way, understanding processes can be seen as operators that search a space consisting of different representational states. These operators generate, modify, and test the adequacy of representations. We see representational states as encompassing more than just what type of diagram is used, additionally they reflect the problem solver's current model of how a task works. Thus we term this space *model space*.

Model Space. In our hierarchical three-space theory of problem solving we propose that model space provides not only the representation of instances, but also defines the rule space to be searched. Which rules appear plausible will depend on how a problem solver thinks the task works. For example, if each component of a system is thought to be independent, then rules proposing interactions will not be considered. If the model changes, then interactions may become part of states in rule space.

Current utility is the criteria for assessing one's state in model space as there is no final goal state, a better understanding of the task may always be possible. So instead of a test for "solution", there may be tests for the adequacy of a model state, that is, does this model seem to work?

Although it violates our application of the term "problem solving" to tasks completed within a few hours, for expository purposes we will illustrate model space with the debate over competing models of light. Two models of light were proposed: a wave model and a particle model. The hypothesis that a scientist would test depended on which model the scientist believed. The wave model suggested that light is a wave, therefore a relevant hypothesis to test was whether light shows interference patterns. The particle model suggested that light is a particle, therefore a relevant hypothesis to test was whether light exerts pressure. Testing these hypotheses led to movement in the model space for light. Neither model was accepted as completely correct, instead the competing models were synthesized into a model in which light was both a particle and a wave. Although this particular movement in model space was slow, it still had the characteristics of a movement in a problem space. There were clearly defined states (initially two different models states, which expanded in number when the possibility of combinations arose), and there were processes for comparing and moving between states (driven by search of rule space). There were no processes for deciding whether the final goal state in model space had been reached, only utility. The current model of light does not rule out the possibility that a new model may emerge.

During problem solving, movement in model space may occur much faster than did movement in the model space for light. Whenever people are faced with a new task, it is necessary to form a model of that task, the current state of which may need to quickly and often be revised, just as Larkin's (1983) physicists did.

Search in rule space can drive movement in model space. For example, if the rules suggested by a model fail, then eventually the response will be to change the model. If the rules make a false claim or mandate an impossible action, then the problem solver can be said to have reached an *impasse* (Brown & VanLehn, 1980). Such impasses require repair procedures, such as when Larkin's (1983) physics experts changed their schemas when faced with a contradiction. In our terms this is movement in model space.

Success in rule space could also lead to movement in model space. While less likely to result in wholesale change to a model, success can lead to modification of the current model, such as through *elaboration*. Elaborations (see VanLehn, 1989) are assertions about the problem without having any impact on previous assertions. Simply filling in slot values in a schema is a form of elaboration, but so are new statements about the representation of the problem which may arise from the testing of rules through the generation of instances.

The hierarchical three-space model of problem solving is represented diagramatically in Figure 1. In this model, the problem description provides the initial model state, which in turn defines the rules space consisting of all possible rules that the model suggests are plausible. The problem solver's state in rule space defines what are the relevant instances and how they should be represented in order to test rules. Instances are then generated by invoking experiments (i.e., interaction with the world) or from memory. The results of generating instances can be used to modify rules, that is, cause movement in rule space (confirmation can be seen as a form of movement too in that the confidence in the rule state would be enhanced). Repeated failure for the rules in the rule space may lead to modification of the model, either directly (e.g., the failed rules may suggest different types of rules), or by evoking search mechanisms in model space in order to overcome the impasse. For each space, memory provides knowledge that is used by the search processes.

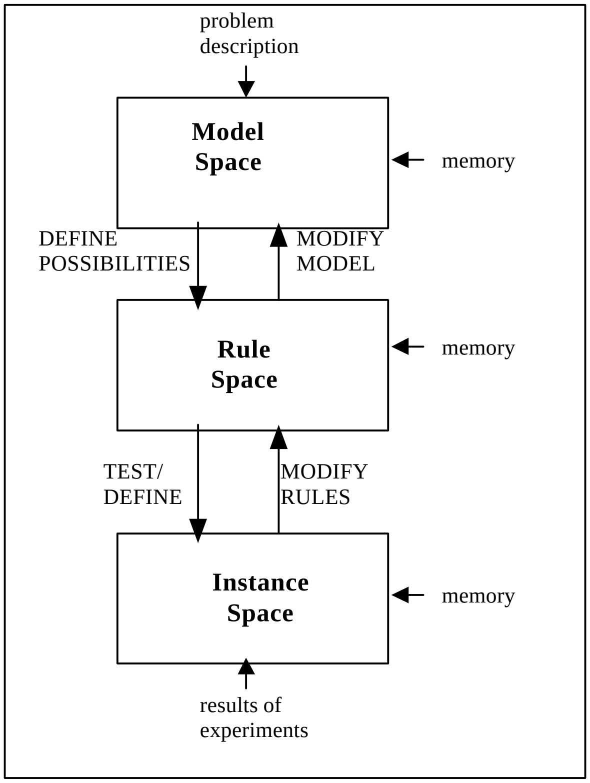

Figure 1. A hierarchical three-space theory of problem solving

We can illustrate these spaces using Vollmeyer et al's (1996) biology-lab task. Initially participants had to construct a model of what sort of system they thought they were faced with. Their model defines a rule space consisting of all possible links between inputs and outputs. If participants hit on the right model immediately, that biology-lab is a straightforward linear system, then they can solve it quickly. To test rules, instances are generated consisting of particular sets of inputs. Such a model defines a constrained, and thus quickly searchable, rule space. However, most participants start out with less precise models. For example, they have models that include the possibility of interactions, or random effects. A model including such possibility defines a larger rule space, and searching these parts of the rule space are at best a waste of time, and at worst confusing.

Search Processes. An advantage of treating problems solving as search of multiple spaces, is that it suggests a series of questions about the nature of the search processes. For each space we have to ask Simon and Lea's (1974) questions: 1) what are the knowledge states, 2) what are the generative processes, 3) what are the test processes, 4) what are the selection processes?

Table 1 is a proposal for the nature of the search processes. Most of the processes invoked are processes already studied. For *instance space* the processes are those normally invoked for problem solving as search of a problem space, but Table 1 also specifies relationships between other processes, such as induction, hypothesis testing, metacognition, and analogy. Table 1 suggests a specific organization between different processes involved in problem solving and learning from problems solving. For *rule space*, induction and hypothesis testing are clearly distinguished as the generative and test processes respectively, and they are both distinguished from analogy. Table 1 also highlights processes we know little about, in particular, the selection processes for rule and model space.

It is clear from Table 1 that we understand least about model space. This is not surprising given that it encompasses the questions of "How do we form representations?" Table 1 implies though a useful way to think of research into analogy, a common topic in recent years (see Holyoak & Thagard, 1995). If analogies give people a new way of looking at a situation or problem (e.g., the water analogy for electricity) then they can be seen as a generative process in model space. Analogies of this type are therefore distinguished from induction.

Table 1: The four generalized problem space components for each of the three spaces.

	Instance space	**Rule space**	**Model space**
knowledge states	states of a task	hypothesized rules	possible models of the tasks
generative processes	operators for changing the state of the task	operators for generating rules (e.g., induction)	operators for generating new models (e.g., analogy)
test processes	evaluate how close current state of task is to its goal state	hypothesis testing (e.g., generate critical instance)	evaluate how well current model fits (e.g., metacognitive processes)
selection processes	select operator or evaluation method	decide which rule to test, or how to generate a rule	select method for evaluating or generating new model

Formalising a Hierarchy of Spaces. In Figure 1, the hierarchical nature of the three-space theory is made clear. We aimed to create a hierarchy because it makes the spaces clearly distinct. We agree with the proposed constraint of Baker and Dunbar (1996) that in multispace theories the spaces at different levels of abstraction (e.g., rule and instance space) should be isomorphic, whereas those at the same level (e.g., different representational forms of the same problem) should be homomorphic. Figure 1 presents the spaces as hierarchical, and we can describe them as being hierarchical, but to truly impose this constraint we need to propose a formal definition that is hierarchical.

To give the spaces a formal definition, we start with the claim that any task can be seen as defined by a set of inputs, a set of outputs, and a set of rules relating those inputs to the outputs. Productions can have this form, so the generality of this claim is wide. In this formalism each output can be seen as a function of the inputs and constants associated with the inputs. Thus, a task with a set of X_1 to X_M inputs and Y_1 to Y_N outputs can be defined by the following set of very general functions:

$$Y_1 = f_1(c_{10}, c_{11}, X_{11}, c_{12}, X_{12}, c_{13}, X_{13}, .., c_{1M}, X_{1M})$$
$$Y_2 = f_2(c_{20}, c_{21}, X_{21}, c_{22}, X_{22}, c_{23}, X_{23}, .., c_{2M}, X_{2M})$$
$$...$$
$$Y_N = f_N(c_{N0}, c_{N1}, X_{N1}, c_{N2}, X_{N2}, c_{N3}, X_{N3}, .., c_{NM}, X_{NM})$$

The relationship between different hierarchical spaces can be specified in terms of the different components of these equations that a state in each space will specify. A *model state* specifies a set of functions with constants left unspecified; a *rule state* specifies a set of constants; whereas a particular set of X values (with resulting Y's) represent *instance states*. For example, consider a task that could be described as a single output with two inputs. This would be defined by a single equation: $Y = f(c_0, c_1, X_1, c_2, X_2)$. A model suggesting that inputs are additive specifies the equation $Y = c_0 + c_1X_1 + c_2X_2$. The rule that "$X_1$ has twice the effect of X_2 but there is no constant effect," is expressed by the equation: $Y = 2X_1 + X_2$. This hypothesized rule could be tested by generating an instance with values of 5 for X_1 and 5 for X_2 and testing if the resulting value of Y is 15. Biology-lab fits easily into this framework as X values can be seen as changes to inputs, constants define particular possible links, and the shape of the functions are the nature of possible rules. However our argument is that any task could be seen in these terms, so applying the hierarchy constraint when determining the exact nature of the spaces for a task can be seen as requiring a specification of how the task fits into this formalism. The mathematics of this formalism are not in themselves insightful, but fitting spaces to this formalism creates constraints on the definitions of the search spaces.

Comparison to Other Approaches

Other Multisapce Models. Ours is not the only work on multispace models. Another is the four-space model of Schunn and Klahr (1996) for scientific discovery. This model differs in various ways from the hierarchical three-space theory, but there is not space here to fully explore the differences. An important difference is that the four-spaces are not constrained to be hierarchical. The scope of the four-space model is not clear, but if it can be a general model of problem solving, then we would welcome it as another attempt to address the lack of theory in problem solving research. Dialectic progress requires competing alternatives.

How do multispace models in general relate to Soar (Newell, Rosenbloom, & Laird, 1989) and ACT-R (Anderson, 1993)? Soar and ACT-R are frameworks in which detailed models of problem solving can be built. Because Soar constructs a new problem space whenever the need arises, Newell (1989) proposed that Soar could model Klahr and Dunbar's (1988) theory, so by extension it can model all multispace theories. Whereas it should be possible to build multispace models in the Soar architecture, they are not equivalent just because they both involve multiple problem spaces. The spaces in multispace models are conceptually distinct and interact in specified ways, so a compatible model built in Soar would have to incorporate these assumptions.

Anderson's (1993) ACT-R does not explicitly incorporate the idea of interactive search of multiple spaces, but there appears to be no reason why it could not model such processes. The current goal in ACT-R is critical, because subgoals encourage the firing of certain sets of productions. Such sets could be considered to define different spaces, so perhaps rapid transition between different subgoals could simulate an interactive search between spaces. The implications of such an approach are unclear.

Situated Cognition. We started by decrying the lack of alternatives theories in problem solving research, but there exists an approach to problem solving that does not focus on search: situated cognition. Situated cognition places a great emphasis on the context of cognition and denies (or at least de-emphasizes) that symbolic processing (such as search of a problem space) lies at the heart of cognition. The extent to which situated cognition is an antithesis to problems solving as search, is not clear. Vera and Simon (1993) tried to place situated cognition into the symbolic framework, but the replies to their article suggested that researchers taking the situated cognition approach see it as fundamentally different. However the problem with situated cognition emerging as an antithesis to the thesis of problem solving as search may be that neither the thesis nor the antithesis is clear enough to begin with.

Like any clearly stated antithesis we would welcome the emergence of a competitor such as situation cognition. Within the three-space model, in general we could try to explain the phenomena that cognition is often heavily context dependent as the claim that movement in the model space is difficult, and may usually define only a restricted rule space. Perhaps this is the general condition, and the implications of this would have to be worked out.

Conclusions

We have argued that a hierarchical three-space theory of problem solving can be derived from existing studies and ideas about problem solving. In constructing this theory we have been guided by Schunn and Klahr's (1996) three criteria for when to propose additional problem spaces in multipspace search theories. The first criterion is logical, do the spaces involve search of different goals and entities? (We would also add, do they use different operators for search?) The three spaces we propose clearly involve different kinds of states, goals, and ways of searching that space, so we think we meet this criterion. The second criterion was, do the spaces differ empirically? There is evidence from existing literature that different factors influence behavior, so we think we can meet this criterion. Schunn and Klahr's third criterion was implementational, spaces should be able to be represented distinctly in a computational model that can perform the task. At the moment we can do more than suggest how such a model using the our theory would work, but constructing such a model is an important aim.

To test the hierarchical three-space theory we intend to examine the testable implications it has for how people may best learn from encountering novel tasks. It suggests that whether hypothesis testing will be a good strategy for learners depends on the quality of the learner's model. Learners with a poor model may be disadvantaged by being encouraged to test hypotheses. A current weakness with the theory is that we may be able to define relatively what are good and poor models in terms of some metric of the size of the rule space the model state defines, but it may be hard to define absolute model goodness. Specifying the distinction between good and poor models precisely is an important aim of future research, especially if we are to investigate the practical implications of the theory. Also required is further study into the reality and properties of the links between spaces proposed by the theory.

Have we met our aim of proposing a theory for problem solving? We are trying to develop the hierarchical three-space theory so that it can generate predictions in terms of the interactions between different spaces, and hope to make the theory a tool for organizing the different processes involved in problem solving. However, we recognize that the theory requires more development, both computationally and empirically, before it is truly more than a framework. Such attempts by problem solving researchers are necessary though because until there are such theories, problem solving will remain just a set of diverse and sometimes unrelated phenomena.

References

Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.

Anderson, J. R. (2000). *Cognitive psychology and its implications* (5th ed.). New York: Worth.

Baker, L. M., & Dunbar, K. (1996, July). *Problem spaces in real-world science: What are they and how do scientists search them?* Paper presented at the Eighteenth Annual Conference of the Cognitive Science Society, San Diego.

Bransford, J. D., & Franks, J. J. (1971). The abstraction of linguistics ideas. *Cognitive Psychology, 2,* 331-350.

Brown, J. S., & VanLehn, K. (1980). Repair theory: A generative theory of bugs in procedural skills. *Cognitive Science, 4*, 379-426.

Burns, B. D. (1996). Meta-analogical transfer: Transfer between episodes of analogical reasoning. *Journal of Experimental Psychology: Learning, Memory and Cognition, 22,* 1032-1048.

Dunbar, K. (1993). Concept discovery in a scientific domain. *Cognitive Science*, *17*, 397-434.

Hayes, J. R., & Simon, H. A. (1974). Understanding written problem instructions. In L. W. Gregg (Ed.), *Knowledge and cognition*. Hillsdale, NJ: Erlbaum.

Hayes J. R., & Simon, H. A. (1977). Psychological differences among problem isomorphs. In N. J. Castellan, D. B. Pisoni, & G. R. Potts (Eds*.), Cognitive Theory, vol. 2*. Hillsdale, NJ: Erlbaum.

Hegel, G. W. F. (1807/1931). *The phenomenology of mind* (J.B. Baille, Trans.). London: Allen & Unwin.

Hofstadter, D. R. (1995). *Fluid concepts and creative analogies*. New York: Basic Books.

Holyoak, K. J., & Thagard, P. (1995). *Mental leaps*. Cambridge, MA: MIT Press.

Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science*, *12,* 1-55.

Larkin, J. H. (1983). The role of problem representation in physics. In D. Gentner & A. L. Stevens (Eds.), *Mental Models*. Hillsdale, NJ: Erlbaum.

Maier, M. R. F. (1930). Reasoning in humans. I. On direction*. Journal of Comparative Psychology, 10*, 115-143.

Maier, N. R. F. (1940). The behavior mechanisms concerned with problem solving. *Psychological Review*, *47*, 43-58.

Newell, A. (1989). How it all got put together. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing*. Hillsdale, NJ: Erlbaum.

Newell, A., Rosenbloom, P. S., & Laird, J. E. (1989). Symbolic architectures for cognition. In M. I. Posner (Ed.), *Foundations of cognitive science*. Cambridge, MA: MIT Press.

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

Schunn, C. D., & Klahr, D. (1996). The problem of problem spaces: When and how to go beyond a 2-space model of scientific discovery. In G. W. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 25-26). Hillsdale, NJ: Erlbaum.

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

Sternberg, R. J. (1995). *In search of the human mind*. Orlando: Harcourt Brace.

VanLehn, K. (1989). Problem solving and cognitive skill. In M. I. Posner (Ed.), *Foundations of cognitive science*. Cambridge, MA: MIT Press.

Vera, A., & Simon, H. A. (1993). Situated action: A symbolic interpretation. *Cognitive Science, 17,* 7-48.

Vollmeyer, R., Burns, B. D., & Holyoak, K. J. (1996). The impact of goal specificity and systematicity of strategies on the acquisition of problem structure. *Cognitive Science, 20,* 75 - 100.

Weisberg, R. W., & Alba, J. W. (1980). An examination of the alleged role of 'fixation' in the solution of several 'insight' problems. *Journal of Experimental Psychology: General, 110*, 169-192.