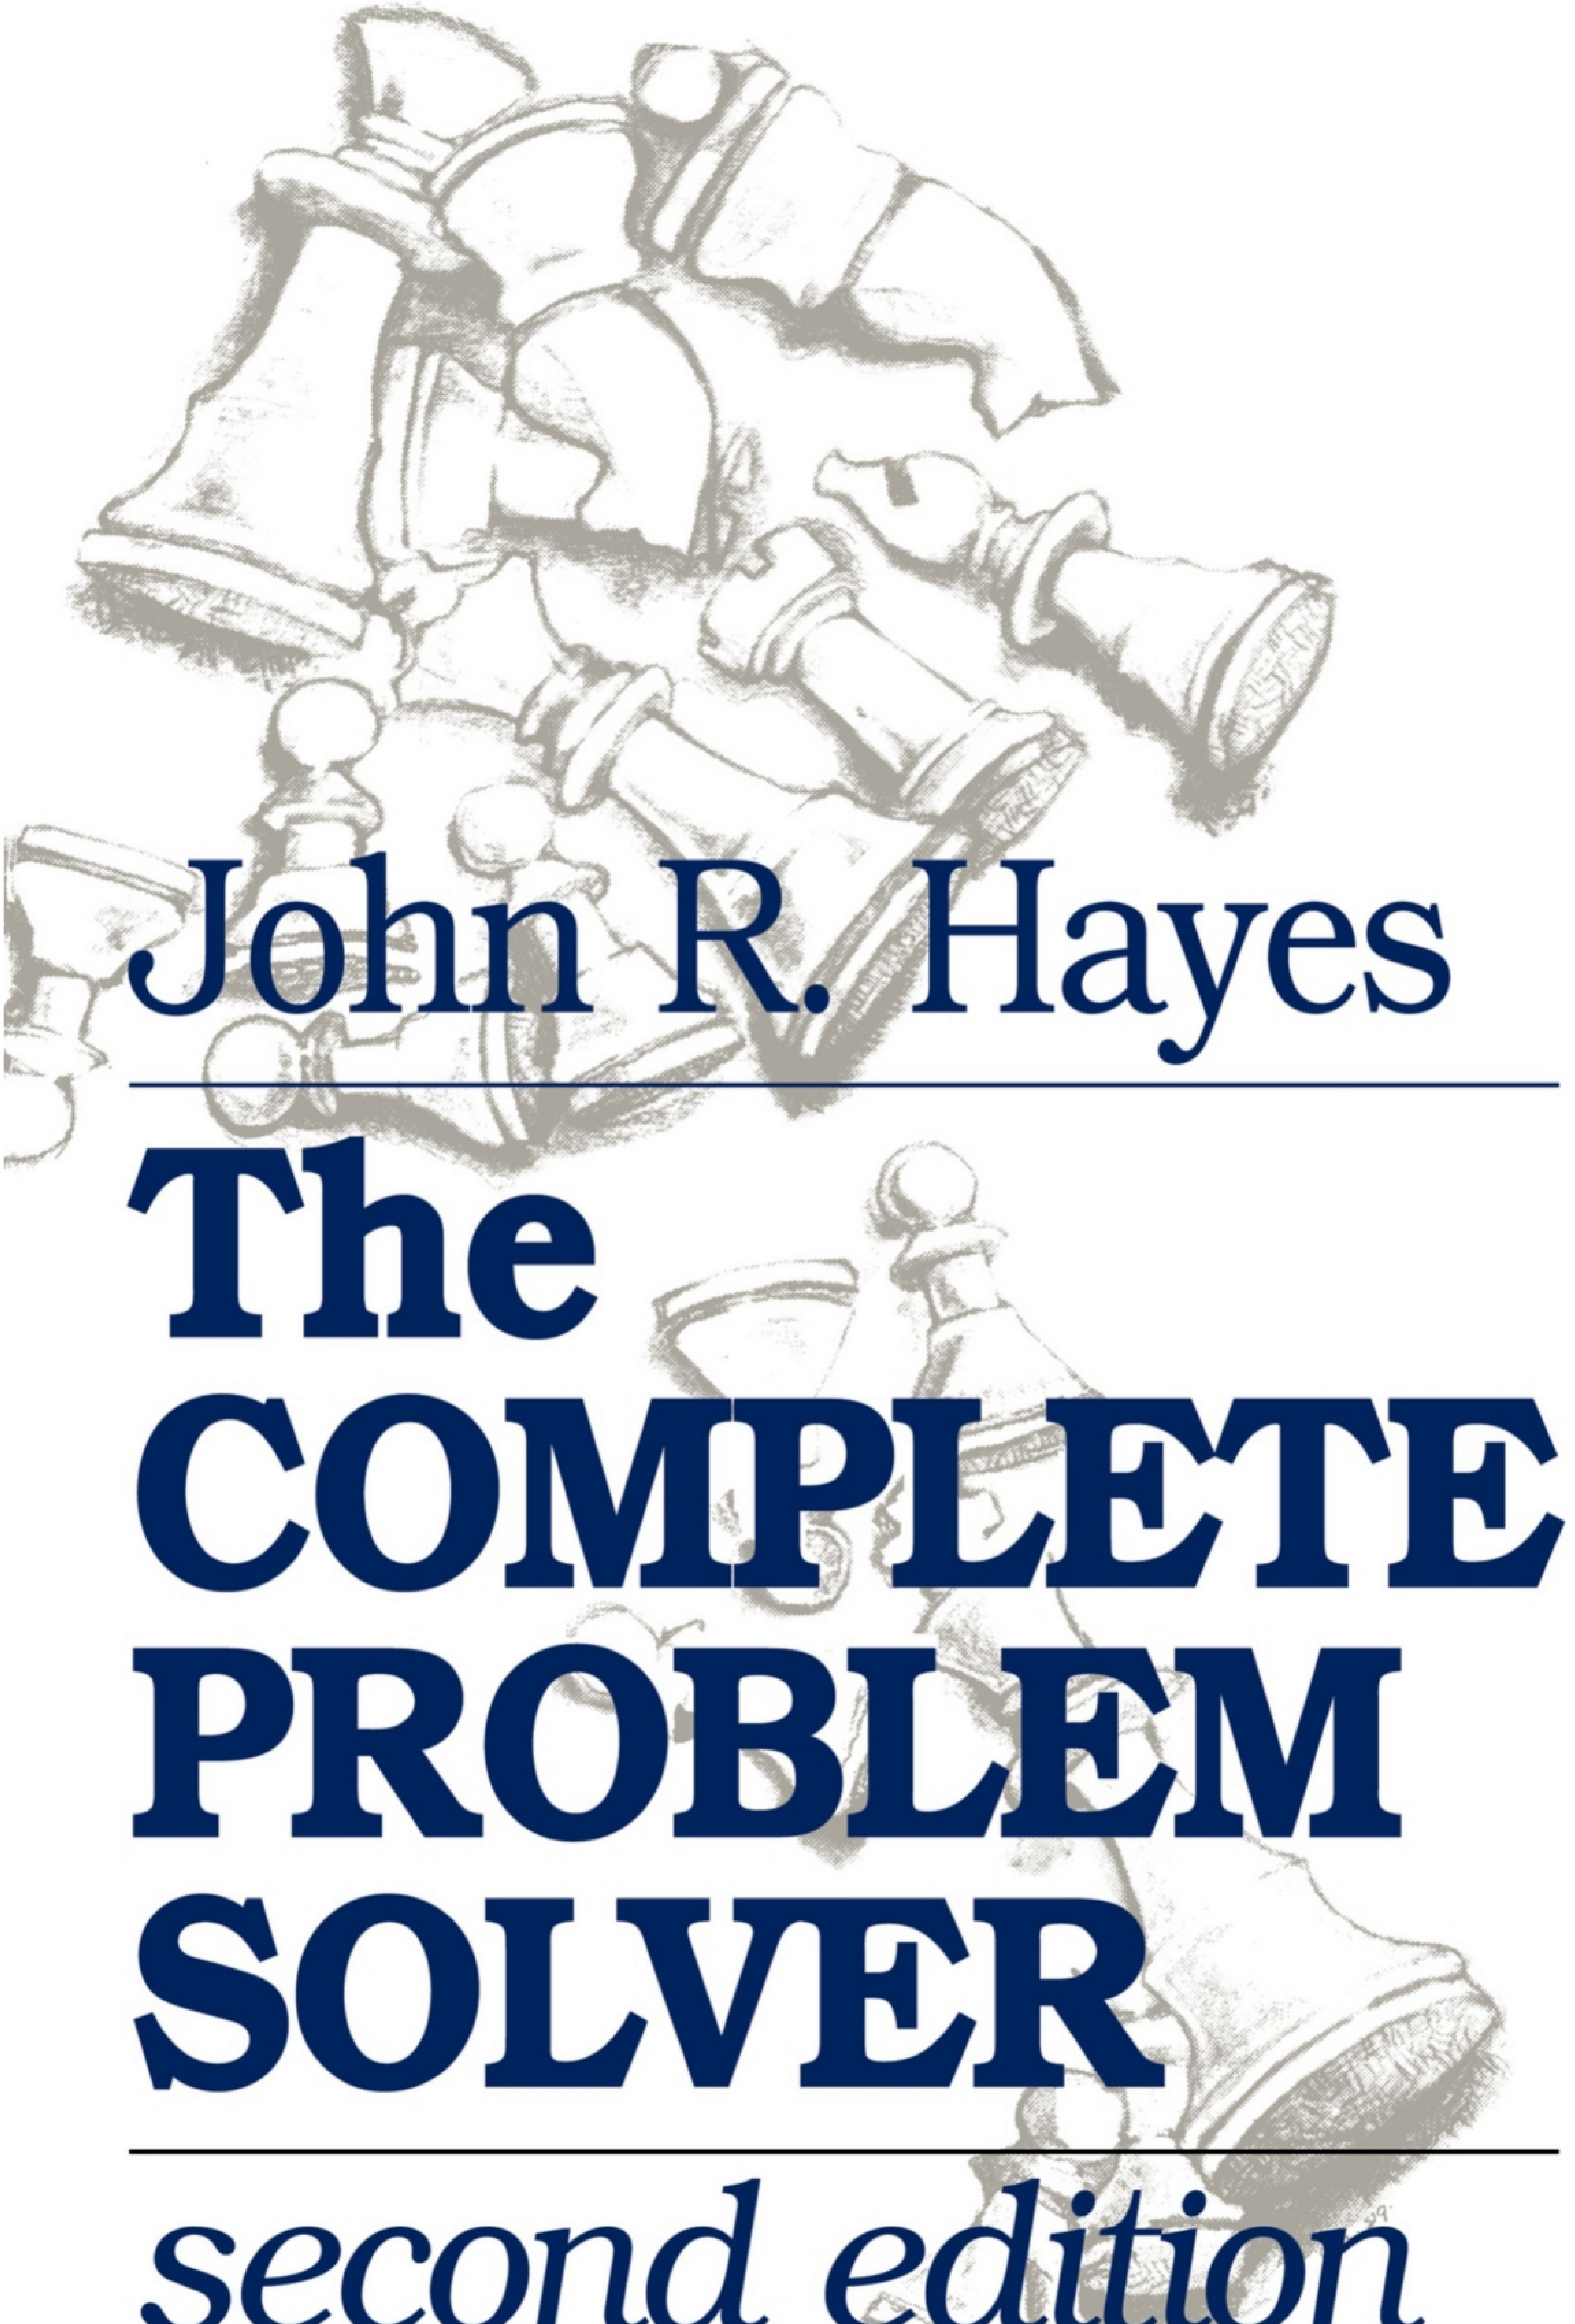

John R. Hayes

The COMPLETE PROBLEM SOLVER

second edition

The Complete Problem Solver

second edition

John R. Hayes
Department of Psychology
Carnegie-Mellon University

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

First Published by
Lawrence Erlbaum Associates, Inc., Publishers
365 Broadway
Hillsdale, New Jersey 07642

Transferred to Digital Printing 2009

by Routledge
270 Madison Ave, New York NY 10016 2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Library of Congress Cataging-in-Publication Data

Hayes, John R., 1940-
The complete problem solver.
Includes bibliographies and index.
1. Problem solving. 2. Thought and thinking.
I. Title.
BF441.H33 1989 153.4'3 88-7105.
ISBN 0-89859-782-X
ISBN 0-8058-0309-2 (pbk.)

Publisher's Note

The publisher has gone to great lengths to ensure the quality of this reprint but points out that some imperfections in the original may be apparent.

Contents

Acknowledgments

This book is the result of many years searching for a satisfactory way to teach a course on general problem solving skills. By general problem solving skills I mean skills that can be used by anyone in solving problems that occur in everyday life.

The idea that I should try to teach a problem solving course arose in a conversation with Herb Simon and Steve Rosenberg in 1974. The three of us decided that if we really knew something useful about problem solving, we ought to be able to teach it. Steve and I, accepting the challenge, put together an initial version of the course which we then taught to a mixed group of engineers and fine arts students at Carnegie-Mellon University. Our first version had many faults, but student response was still positive. I was sufficiently encouraged to continue. The course then underwent a long evolution as I borrowed ideas from a wide variety of sources and tested them in the classroom. Many of these ideas came from people who, at one time or another, have taught the course with me:

- Steven Rosenberg—who helped me to start the whole enterprise,
- Linda Flower—who made me see the relevance of problem solving to the crucial skill of writing,
- Lee Gregg—who encouraged me to think big about the course, and
- Lynne Reder—who added her sophistication in the fields of memory and decision making.

Other ideas came from people who were generous enough to contribute guest lectures. Delarese Ambrose spoke about how to get along in small groups; Edward Constant, about cultural influences in creativity; John Gatchnig, about recursive methods in problem solving; Douglas Lenat, about heuristics; David Meeker, about information retrieval; Robert Neches, about listening to lectures; John Payne, about decision making; Herbert Simon, about series completion problems; and Richard Teare, about problem solving in engineering.

I would like to add a special note to thanks to John Payne. It was through his example and encouragement that I added the section on decision making to the course—a section which many students have found extremely useful.

In teaching a course on problem solving skills, it is extremely important to provide students with close personal supervision. Teaching problem solving skills is a bit like coaching: The instructor needs to watch the students in action to be sure that they are performing the skills in the right way. As course enrollment increased, it became impossible for me to give this supervision myself. Instead, I relied on a group of teaching assistants—mostly undergraduates—who volunteered to supervise groups of about 15 students in weekly sections and to meet weekly with me to discuss section problems. The success of the course has depended very heavily on the efforts of the teaching assistants listed here:

Doug Bauman
Cynthia Berkowitz
Sandra J. Bond
Barbara Madera Clevenson
Dana Dunn
Suzanne Eckert
Carole Elm
Anna-Lena Ericsson-Neches
Becky Freeland
Amy Gift
Richard Gorelick
Mark Hanna
Ron Kander
Anne Karcher
Rich Kleinhample
Jeanne Kravanja
Jamie Leach
Kevin Brown
Steve Ciampi
Aaron Clevenson
Anne Lux
Bill Lyden
Marilyn Mantei
John Maslany
David Meeker
Ernie Prescott
Greg Pisocky
Susan Robinson
Mark Segal
Lisa Thaviu
Judi Vitale
Philip Werner
Ellen Zoll

Jeanne Halpin has taught the course very successfully on her own. By doing so, she has demonstrated the important point that the course is genuinely "exportable."

Finally, I want to note the very important contributions of Sandra Bond. She has coordinated the complex mechanisms of the course: making sure that lectures got delivered, that grading got done, that TA's with difficulties got listened to, and that hundreds of students found their way to weekly section meetings. She served as a teaching assistant for many years, consistently receiving excellent evaluations from her students. In addition, she guest lectured on creativity in women. She has also contributed

heavily in the preparation of this book. She typed the manuscript, drew the figures, edited and proofread the text, and did much of the research and writing for the final chapter. In short, a great deal of the work you see in this book is hers.

In appreciation, I dedicate this book to the many mentioned above who have contributed to this project.

Introduction

This book has two purposes. It is designed to provide you with skills that will make you a better problem solver, and to give you up-to-date information about the psychology of problem solving.

The first purpose is clearly a practical one, but I believe the second purpose is, too. It is important for people to know how their minds work. Certainly for humanistic reasons—knowledge of our human nature is valuable in itself—but it is also important because it provides us with a degree of flexibility which we might not otherwise have. If we can examine or own problem solving processes with some degree of understanding, then we have a better chance of improving them. Further, if we have some understanding of how people think, we can be more effective in helping others. Anyone who is to teach, or to tutor, or even to help a child with homework, can benefit from knowledge of how human problem solving processes work and how they can go wrong.

Early in my career as a psychologist, a student asked me about my special area of interest. I told him that I studied people's thinking processes. "Oh, thinking!" he said, "I know all about that. I'm a math major." Of course, he did know a lot about thinking—he knew about *how to do it,* at least in certain cases. Given a math problem, he could draw on a wealth of experience to help him find a solution. But if he were like most people, he would have a very difficult time articulating that wealth of experience; he knew how to think but he didn't know how to describe his own thinking. When they are faced with their first teaching task, whether in school or out, many professionals discover a vast difference between their ability to do what they do very well and their ability to describe what they do to others.

In this book, then, we hope to provide you with some skills that will help you to solve problems, but we also hope to provide you with some knowledge that will give you greater insight into what you are doing and an increased ability to understand others.

What Is a Problem?

If you are on one side of a river and you want to get to the other side but you don't know how, you have a problem. If you are assembling a mail-order purchase, and the instructions leave you completely baffled about how to "put tab A in slot B," you have a problem. If you are writing a letter and you just can't find the polite way to say, "No, we don't want you to come and stay for a month," you have a problem. Whenever there is a gap between where you are now and where you want to be, and you don't know how to find a way to cross that gap, you have a problem.

Solving a problem means finding an appropriate way to cross a gap. The process of finding a solution has two major parts: *1*. Representing the gap—that is, understanding the nature of the problem, and 2. Searching for a means to cross it.

Representing the Gap

If people fail to understand the nature of the gap, they may well set off in the wrong direction to search for the solution. Suppose you told a friend that you would give him $10,000 if he put his elbow in his ear. "Easy," your friend says; "I'll just cut off my elbow and put it in any ear you choose." Now you may question your friend's values, but you are also pretty sure that he understands the nature of the difficulty—the gap—that the problem presents. On the other hand, if your friend said, "Easy, I'll stand on a chair," you would suspect that he didn't really understand the nature of the difficulty.

Representing the gap isn't always easy. In fact, the main difficulty in many problems is just the difficulty of representing the gap. Consider the Driver's License Problem.

Problem 1. The Driver's License

> When Tom and Bill applied for their drivers' licenses, they were asked their ages. Bill, who was a bit of a revolutionary, said they were both in their twenties and that was all he was going to reveal to a bunch of bureaucrats. The clerk insisted on more specific information so, to smooth things over, Tom added that they both had the same birthday, and that he was four times as old as Bill was when he was three times as old as Bill was when he was twice as old as Bill was. At this the clerk fainted and the two snatched up their licenses and disappeared. When the clerk came to and realized that he would have to complete his records some way or other he began to do a little figuring, and before long had found out how old the two were. Can you tell, too?

A typical reaction to this problem is to say, "What?" or beat a hasty retreat explaining, "I never was much good at puzzles." But the problem really isn't very difficult once we find an appropriate representation for it. In Chapter 1, we will discuss processes by which we come to understand the nature of a problem; we will show that the way we represent the gap can make an enormous difference in the difficulty of the problem; and we will provide some hints on how to represent problems to make them easier.

Finding a Solution Path

Once we understand the nature of a problem, there are still many reasons why we may have difficulty in finding a solution to it. The problems below illustrate some of the most important reasons.

Problem 2. The Loser

> A man once offended a fortune teller by laughing at her predictions and saying that fortune telling was all nonsense. He offended her so much, in fact, that she cast a spell on him which turned him into both a compulsive gambler and in addition a consistent loser. That was pretty mean. We would expect the spell would shortly have turned him into a miserable, impoverished wreck. Instead, he soon married a wealthy businesswoman who took him to the casino every day, gave him money, and smiled happily as he lost it at the roulette table. They lived happily in just this way ever after. Why was the man's wife so happy to see him lose?

The story poses a problem for most of us when we first see it. It would be no problem if the man were winning money. We know right away how to get from winning money to happy smiles, but to get from losing money to happy smiles *is* a problem—there is a gap that we can't immediately cross. How can losing money lead to happiness?

In trying to bridge the gap, people propose a variety of solutions:

"Perhaps she is so rich that she really doesn't care about the money."

"Perhaps she is becoming a nun and wants to give all her money away."

"Perhaps her crazy grandfather left a will which required her to lose all her money by 21 in order to inherit a billion,"

"Perhaps she is a masochist."

These solutions vary in quality. The solution about the woman becoming a nun has the difficulty that it ignores her husband. A solution which seems to us better than all of these is this: When playing roulette, the man bets, say, on red and loses, as his spell requires. The woman, however, bets twice as much on black and wins. In short, she has turned her husband's misfortune into an advantage. His loss is their gain, and so the smiles. The gap is crossed.

The problem illustrates a very important process in problem solving—the process of invention. In many problems, lots of approaches are conceivable-some of them better than others. Typically, a person will try several approaches before hitting on a good one. If people can't think up any approaches, then they can't solve the problem.

Invention is an important problem solving process, but it isn't the only process required in solving problems. There are many problems in which invention is easy but the problem is still difficult.

Problem 3. The Combination Lock

Suppose that you have the problem of opening the 10-dial combination lock shown above. Proposing possible solutions is easy. The dial setting shown may be a solution—and it may not. There is a total of 10^{10} or 10,000,000,000 or ten billion possible dial settings, any *one* of which may be the solution. This is where the difficulty of the problem lies—finding the single correct combination among ten billion possibilities. If we tried one combination every second, working day and night, it would take us 317 years to try them all. Some problems, then, like the Combination Lock Problem, are difficult to solve because we have to search for the solution among a very large number of alternatives.

The following problem is a difficult one, even though it involves neither invention nor examining large numbers of alternatives.

Problem 4. The Rational Investor

> Suppose that you have a choice between a safe investment which yields a sure 25% return and a risky investment which gives you an even chance of either tripling your money or of losing it. Which investment is best?

The difficulty of practical decision making demanded by problems like this one lies in evaluating the alternatives. People are often unreliable when making such evaluations. If they choose from the same alternatives on several different occasions, the results may be quite inconsistent.

Which of the two investments is best depends in a complex way on the investor's financial circumstances. To evaluate the alternatives accurately, most people require explicit decision procedures such as those described in Chapter 9.

Some problems are difficult because we have trouble remembering where we are on our path to the solution. Try to solve problem #5 before reading further.

Problem 5. Cats Among Pigeons

> Messrs. Downs, Heath, Field, Forest, and Marsh—five elderly pigeon fanciers-were worried by the depredations of marauding cats owned by five not less elderly ladies, and, hoping to get control of the cats, they married the cat owners.
>
> The scheme worked well for each of them so far as his own cat and pigeon were concerned; but it was not long before each cat had claimed a victim and each fancier had lost his favorite pigeon,
>
> Mrs. Downs'cat killed the pigeon owned by the man who married the owner of the cat which killed Mr. Marsh's pigeon. Mr. Downs' pigeon was killed by Mrs. Heath's cat. Mr. Forest's pigeon was killed by the cat owned by the lady who married the man whose pigeon was killed by Mrs. Field's cat.
>
> Who was the owner of the pigeon killed by Mrs. Forest's eat?
>
> (from Phillips, 1961)

Unless you are an expert in solving this sort of problem, you may have had some difficulty in keeping track of your place on your way to the solution. You may have found yourself asking questions like, "Wasn't Mrs. Marsh the lady who ate the cat that married Mr. Forest's pigeon?—Or was it the other way around?" Being able to remember your place on the solution path is a critical problem solving skill.

Consider Problem #6, but don't consider it for very long.

Problem 6. Who's Got the Enthalpy?

> Liquid water at 212°F and 1 atm has an internal energy (on an arbitrary basis) of 180.02 Btu/lb_m. The specific volume of liquid water at these conditions is 0.01672 ft^3/lb_m. What is its enthalpy?
>
> (from Smith and Van Ness, 1959)

Problem *#6* is not a very difficult problem if you know something about thermodynamics. If you don't, however—if, for example, you haven't the foggiest idea what enthalpy is—then it's an impossibly hard problem. I present this problem not to make you feel bad, but to dramatize the extreme importance of knowledge in problem solving. If you are missing relevant knowledge, an easy problem may appear difficult or impossible. If your knowledge of math and science is weak, the problems that scientists solve may appear much harder to you than they really are. If the humanities or the arts are your weak suit, then people who can understand philosophy or who can interpret a musical score may seem magically intelligent to you. The moral is this: Much that passes for cleverness or innate quickness of mind actually depends on specialized knowledge. If you acquire that specialized knowledge, you too may be able to solve hard problems and appear clever to your less learned friends.

Organization of the Book

The six problems given here illustrate six important aspects of human problem solving which we emphasize:

Problem 1: Representation
Problem 2: Invention
Problem 3: Search for the Solution Among Many Alternatives
Problem 4: Decision Making
Problem 5: Memory
Problem 6: Knowledge

The book is divided into four sections:

Section I. Problem Solving Theory and Practice (representing problems and searching for solutions)
Section II. Memory and Knowledge Acquisition
Section III. Decision Making
Section IV. Creativity and Invention

While the order in which you read Sections 1, 2, 3, and 4 is not critical for understanding, I do recommend that you read Section 1 first for an overview of the problem solving process.

References

Phillips, H. *My Best Puzzles in Logic and Reasoning.* New York: Dover, 1961.
Smith, J. M,, and Van Ness, H. C. *Introduction to Chemical Engineering Thermodynamics,* Second Edition. New York: McGraw-Hill, Inc., 1959.

Problem Solutions

Page xii. The Driver's License:
Tom is 24, Bill is 21

Page xv. Cats Among Pigeons:
Mr. Heath

I

Problem Solving Theory and Practice

1
Understanding Problems: The Process of Representation

Usually when we solve a problem, we put most of our attention on the problem and very little attention on ourselves—that is, on what we are *doing* to solve the problem. If we did attend to our own actions, we might notice that they often occur in a characteristic sequence:

1. **Finding the Problem:** recognizing that there is a problem to be solved.
2. **Representing the Problem:** understanding the nature of the gap to be crossed.
3. **Planning the Solution:** choosing a method for crossing the gap.
4. **Carrying Out the Plan**
5. **Evaluating the Solution:** asking "How good is the result?" once the plan is carried out.
6. **Consolidating Gains:** learning from the experience of solving.

This sequence of actions is illustrated in the following problem.

Action	*Problem*
Finding the Problem	I observe Smith, who claims to be too poor to repay the $50.00 he owes me, buying round after round of drinks for his friends.
Representing the Problem	I conclude that Smith is not sufficiently serious about repaying his debt.
Planning the Solution	I consider a polite telephone call or a note remind-

	ing Smith of his indebtedness, but decide instead to ask three very large friends of mine to call on Smith in person.
Carrying Out the Plan	I call my friends, who then deliver my message to Smith.
Evaluating the Solution	Since Smith paid up rapidly without major bloodshed, I regard the problem as satisfactorily solved.
Consolidating Gains	I revise my rules for lending money to Smith and reflect on the value of having a few large friends.

In easy problems, we may go through these actions in order and without any difficulties. In hard problems, though, we may have to do a great deal of backtracking. For example, when we evaluate what we have done, we may decide that our solution is terrible, e.g., "Asbestos bread will not solve the burned toast problem!!" and go back to planning. Or while trying to execute a solution, we may discover something about the problem which will lead us to represent it in an entirely new way—"Oh, now I see what kind of a problem it is!" Retracing of this sort is characteristic of problems that are called "ill-defined." We will discuss these in much more detail later.

Our success as problem solvers depends on the effectiveness with which we can carry out each of the six actions just described. In this chapter, we will examine the nature of problem representations and the processes people use to form them. In addition, we will describe techniques for improving representations so that they make problem solving easier. In the next chapter we will discuss planning, executing, evaluating, and consolidating. We will delay the discussion of problem finding until the final section of the book because this topic is so closely related to the topic of creativity.

How Do People Understand Problems?

Suppose we were to spy on people as they were trying to understand a new problem, such as the Monster Problem below.

Monster Problem #1

Three five-handed extra-terrestrial monsters were holding three crystal globes. Because of the quantum-mechanical peculiarities of their neighborhood, both monsters and globes come in exactly three sizes with no others permitted; small, medium, and large. The medium-sized monster was holding the small globe; the small monster was holding the large globe; and the large monster was holding the medium-sized globe. Since this situation offended their keenly developed sense of symmetry, they proceeded to transfer globes from one monster to another so that each monster would have a globe proportionate to its own size.

Monster etiquette complicated the solution of the problem since it requires: *1.* that only one globe may be transferred at a time, *2.* that if a monster is holding two globes, only the larger of the two may be transferred, and *3.* that a globe may not be transferred to a monster who is holding a larger globe.

By what sequence of transfers could the monsters have solved this problem?

We might see people reading the problem over several tunes and pausing over the hard parts. We might see them drawing sketches or writing symbols on paper, and we might hear them mutter to themselves, something like: "Let's see ... If a monster is holding two globes . . . What does this mean? . . ." If we were to ask people to "think aloud" as they worked on the problem, we would find that their reading, sketching, and muttering reflected a whirlwind of internal activities—imaging, inferencing, decision making, and retrieving of knowledge from memory—activities which are directed toward "understanding the problem." If we look in more detail, we would find that people are selecting information and imaging objects and relations in the problem. For example, after reading the first line of the Monster Problem, a person might form a visual image of three blobs, each touching a circle. The imagined blobs and circles, of course, correspond to the monsters and the globes, and touching in the image corresponds to the relation of holding. The images usually reflect some selection of information, e.g., the blobs may have no hands, or the circles may give no indication that the globes are crystaline.

To understand a problem, then, the problem solver creates (imagines) objects and relations in his head which correspond to objects and relations in the externally presented problem. These internal objects and relations are the problem solver's *internal representation* of the problem. Different people may create different internal representations of the same problem.

Frequently, problem solvers will make an *external representation* of some parts of the problem. They do this by drawing sketches and diagrams or by writing down symbols or equations which correspond to parts of the internal representation. Such external representations can be enormously helpful in solving problems.

The Relation of Internal and External Representations

Sometimes we can solve a problem using only an internal representation. For example, most of us can multiply 17 by 23 entirely in our heads and, with a little effort, get the right answer. Many problems, however, are very difficult to solve without the aid of an external representation. The Monster Problem and the Driver's License Problem in the Introduction are examples of such problems. While it is possible to solve the Monster Problem entirely mentally, it is very difficult to keep track of where you are in this problem without an external representation. You find yourself asking questions like, "Did I give the small globe to the big monster or didn't I?" In the Driver's License Problem, if you don't invent and write down a good algebraic notation, you are very likely to confuse such things as Tom's age *now* with his age at an earlier time.

External representations, then, are often very helpful m solving difficult problems. We should note, though, that external representations *can't help us at all unless we also have an internal representation of the problem.* Imagine that we are playing chess. In front of us the chess board and pieces provide a very useful external representation of the chess game. But when we make a move, we typically try it in our heads before making it on the board. Planning is done internally. Further, we couldn't make moves either in our heads or on the board if we didn't have an internal representation of how each piece moves. In short, intelligent play would be impossible without an internal representation.

In summary:

1. An internal representation is essential for *intelligent* problem solving. Internal representations are the medium in which we think, in the same way that words are the medium in which we talk. Without internal representations, we can't think through the solution of a problem, just as without words we can't speak.
2. Sometimes an internal representation is sufficient for solving. If we were very skillful, we could play "blindfold chess," that is, we could play using only our internal representation, but it wouldn't be easy.
3. For many problems, an external representation is *very* helpful. We will explore how external representations can help later in this chapter.

What Do We Need to Represent in an

Internal Representation?

Consider the Monster Problem discussed previously. If we are to solve this problem, there are four problem parts that we need to include in our internal representation:

1. **The Goal—**where we want the globes to be when we are done.
2. **The Initial State—**that is, which monsters have which globes at the beginning of the problem.
3. **The Operators—**the actions that change one problem state into another— in this case, passing globes back and forth; and
4. **The Restrictions on the Operators—**Monster Problem rules 1,2, and 3.

Here is another problem:

Starting with the arrangement of dots shown below:

Initial State

Try to produce this arrangement:

Final State

Given that a dot can move to an adjacent space on either side, e.g.,

Operator 1

and a dot can jump over *one* other dot of either color into an empty space, e.g.,

Operator 2

However, the white dots can only move to the left and the black dots to the right.

Restriction

Try to identify the goal, the initial state, the operators, and the restrictions in the following problem:

> A farmer traveling to market took three possessions with him; his dog, a chicken, and a sack of grain. On his way, he came to a river which he had to cross. Unfortunately, the only available transportation was an old abandoned boat that would hold only himself and one of his possessions. Taking his possessions across one at a time posed a problem, however. If he left his very reliable dog with the chicken, the dog would very reliably eat the chicken. If he left the chicken with the grain, the chicken would eat the grain and then burst, improving neither of them.
>
> How did the farmer manage to get all his possessions safely across the river?

While all four problem parts are essential in these two problems, this isn't always the case. All problems involve at least a goal, but many problems omit one or more of the other three parts. Suppose a friend says to us, "Get to my house at 10 o'clock." That statement specifies the goal you are to accomplish, but nothing else. It doesn't specify where you should start—north, south, east, or west. There is no special initial state. Further, it doesn't matter how you get there—you can walk, hop, skate, unicycle, take a cab, a helicopter, a large bird, anything—it doesn't matter—no operator is specified. Further still, no restrictions were specified—e.g., "If you hop, use only the left foot," or "If you come by bird, don't use a sparrow." Some other problem statements which specify only a goal are: "Be a success, my child," and, "Prove your point."

Some problems specify just initial state and goal: "Make a silk purse out of a sow's ear"; or initial state, goal, and restrictions: "Make a silk purse out of a sow's ear, but don't smell up the house." Others specify just goal and operators: "Paint a picture," or, "Get to my house by taxi"; or goal, operators, and restrictions: "Drive me home, but don't drive too fast." Finally, we have problems which specify goal and restrictions, e.g., "Build a fire but don't use matches."

To form an adequate internal representation of a problem, we must represent the goal of the problem, and in addition—for problems in which they are required— the initial state, the operators, and the restrictions.

How Internal Representations Are Formed

At first, we might imagine that forming internal representations is a copying process in which the problem solver makes a sort of mental xerox of an external situation—reproducing everything in the external situation and adding nothing. In fact, an internal representation is far from being a copy. Forming a representation is a very active process in which the person adds and subtracts information, and interprets information in the original situation. When you read the Monster Problem, you may have pictured creatures arranged in a row either horizontally or vertically. You may have pictured them in the order they were mentioned in the text—medium, small, large—in order of size, or in some other order. However you pictured them, you *added details* to the representation. The problem said nothing at all about how the monsters were arranged. You may also have added shapes for the monsters such as those shown in Figure 1.

Figure 1. A Representation of the Monster Problem

Selecting Information

While you probably added some details, you probably left others out. For example, your image of the problem situation probably didn't contain anything about "the quantum-mechanical peculiarities" of the monsters' neighborhood. Very likely you regarded this material as "just part of the cover story" and not really relevant to the solution of the problem. Further, you may have recognized that the number of monster hands is irrelevant and left that out of your representation as well. Relevance judgments such as these are useful because they allow us to pare our representations down to manageable size.

In a study exploring relevance judgments in problem solving (Hayes, Waterman, and Robinson, 1977), the experimenter read problems aloud which the subjects had not heard before and asked them to make relevance judgments. The problems were presented in small pieces so that the subjects could make separate judgments about each piece. A typical session for the "Allsports Problem" proceeded as follows:

Experimenter (reading problem): "I went to tea."

Subject:	"Not relevant."
Exp:	"Yesterday."
Sub:	"Relevant—this may be some sort of a time problem"
Exp:	"With an old friend."
Sub:	"Maybe the 'old' is important if this is a time problem. No, I'll say irrelevant."
Exp:	"Mrs. Allsports."
Sub:	"Relevant. Probably an important person."
Exp:	"She has three daughters: Amelia, Bella, and Celia."
Sub:	"OK, now we're into it. It's going to be about the daughters. Relevant."
Exp:	"On the doorstep, I met another friend."
Sub:	"Irrelevant."

Even on first reading, subjects were quite accurate in their relevance judgments. They correctly identified more than 80 percent of the material in the text which was actually relevant, while rejecting more than 40 percent of the text as irrelevant. Clearly, relevance judgments can help us to focus on the important parts of the problem and thus make our task of building a representation easier.

How do people make relevance judgments? The Hayes, Waterman, and Robinson study suggests that one very important factor is the person's knowledge of problem types. A sophisticated problem solver recognizes many problem types, such as distance-rate-time problems, age problems, river-crossing problems, and so on. Once the problem solver has identified the problem type, judging what is relevant for solution is much easier. For example, in another study (Hinsley, Hayes, and Simon, 1977), the experimenter had presented only the first three words of the problem, "A river steamer . . ." At this point the subject said, "It's going to be one of those river things with upstream, downstream, and still water. [She was right.] You are going to compare times upstream and downstream—or if the time is constant it will be distance." By recognizing the problem type, the subject could predict that any mention of speed either upstream or downstream would be important for solution, and that the names of the boat, the captain, and their desination would be irrelevant.

If you want to be quick in finding the essential parts of a problem, knowledge of problem types can be a great help.

Using Knowledge to Interpret Problems

When we form a problem representation, we not only add information and delete information, we also interpret information —that is, we use our knowledge of the language and the world to understand problem information. Imagine the following scene: A room in which we find a father, a mother, a son, and a baby. The father says, "Pedro, Juarnta is crying. Please change her." We infer that the father is talking to the boy, because we know Pedro is a male name and because the comment isn't an appropriate one to address to a baby. We infer from our knowledge of people's behavior at various ages that it is the baby who is crying and in need of a change, not the mother. Further, we infer that the baby is female and named Juanita, and that the family is probably Spanish or Mexican. Finally, our knowledge of family life suggests that when the father says, "Please change her," he does not mean, "Turn her in for a new model."

Even a simple situation like this one requires us to make a number of knowledge-based inferences in order to understand it. We make these inferences so naturally and automatically that often we believe that our conclusions were actually spelled out in the problem and that they were not inferences at all. For example, if we were to ask someone, "How do you know the boy's name was Pedro?" the response might well be, "It said so in the problem." It didn't. That was an inference based on the problem solver's knowledge.

A time when we become acutely aware that we need knowledge to interpret problems is when we don't have that knowledge. When you read through the Enthalpy Problem in the introduction, you may have found yourself asking questions such as, "What is internal energy?" "What's a BTU?" "And what in the world is enthalpy?" Without knowledge of concepts and relations in thermodynamics to build on, we can't represent the problem. We don't understand the initial state or the goal, and we have no idea what the operators might be.

Analogies and Schemas

Very often, when we encounter a problem, we recognize that we have seen a similar problem before. For example, suppose that while you are driving, your car begins to lose power. It will only make 30 miles per hour on level road and slows to a crawl up hills. Also it has a terrible tendency to stall when the traffic light turns green. The first time this happens, you may think that your car is about to die a ghastly death. If you have been through it before, though, you may recognize the symptoms of a familiar problem. Your car (you hope) just needs a tune up. With luck, new spark plugs and points will make it healthy and happy again.

Ahmed and George

For another example, consider the following problem. Two ancient Egyptians, Ahmed and George, were measuring a field on the banks of the Nile. Starting from one corner of the field, Ahmed walked 20 cubits south and George walked 60 cubits west. How far apart were they at this point?

Now, I'm reasonably sure that you have never heard this problem before, since I just made it up. Nevertheless, you likely recognize that this is a "triangle" problem, and because you have a triangle problem *schema*, you know that you should use the Pythagorean theorem to solve it.

There are many familiar problem schemas. For example, there are schemas for distance-rate-time problems, triangle problems, interest problems, river-current problems, river-crossing problems, mixture problems, age problems, and many more. A problem schema is a package of information about the properties of a particular problem type. A schema for triangle problems, for example, may include information that:

1. The initial state will specify lengths of some of the sides of a right-angle triangle;

2. The goal will be to find the length of another side; and
3. The operator will involve application of the Pythagorean theorem.

There is a variety of "optimist" story which (inadvertently) illustrates the importance of our knowledge of problem schemas in representing problems.

Optimist Story 1

> An optimist put a new kind of furnace in his house and found that it cut his heating bills in half. Delighted, he had another one installed, expecting that he would cut his fuel bill to zero.

Optimist Story 2

> An optimist really likes his doctor except that every time he visits his office, he has to wait an hour to see him. Then a brilliant idea strikes him. He decides that if he takes two friends with him to help, he should only have to wait for 20 minutes.

Now, the peculiar thing about the optimists' thinking is not that they are failing to use knowledge, but rather that they are using the knowledge inappropriately. There are many situations in which it is true that if one of something does half a job, then two of them will do the whole job. If one can of paint covers half of the house, then two cans ought to cover the whole house. The optimist's error is that he has applied this schema to heating houses, where it is not appropriate.

The optimist in the second story uses a schema which is perfectly appropriate in "work" problems. If one person can do a job in an hour, three people ought to be able to do the job in 20 minutes. However, there are many activities that can't be hastened by having several people combine their effort. These include waiting, falling off cliffs, and maturing—If one boy reaches puberty at 12, could 12 boys reach puberty at one?

Problem schemas are an important part of the knowledge we use to solve problems. However, as the optimist stories show, we also need to know when the schemas are appropriate and when they are not.

Individual Differences in Problem Representation

Even when two people represent the same problem, they may well not represent it in the same way. A person who is very good at filtering out irrelevant detail may produce a very spare representation, as in Figure 2. Another person who is not good at filtering out irrelevant detail may produce a complex and ornate representation, as in Figure 3.

There are more differences between representations, though, than just the amount of detail they contain. One person may represent a problem in visual imagery, another in sentences, and a third in auditory images. If two people represent a problem in visual images, they may well not use the same images. For example, in imagining the monsters in the Monster Problem, some saw them arranged horizontally, some vertically, and some in a circle.

Figure 2. A Spare Representation

Figure 3. An Ornate Representation of the Same Problem

Our skill in problem solving depends in a very important way on our store of problem schemas. Each problem schema we know gives us a very valuable advantage in solving a whole class of problems—an advantage which may consist in knowing what to pay attention to, or how to represent the problem, or how to search for a solution, or all three. Clearly the more schemas we know, the better prepared we are as problem solvers.

While our problem skill depends on how many sehemas we have, it also depends on the nature of those sehemas, McDermott and Larkin (1978) have shown that novices in physics are more likely to have sehemas that are tied to concrete aspects of the problem situation, e.g., "spring problem" sehemas and "balance problem" sehemas, whereas experts are more likely to have sehemas tied to abstract physics principles, e.g., "energy" sehemas and "moment of inertia" sehemas.

In the same way, inexperienced math students are likely to use separate schemas for the following problems:

Mr. Lloyd and Mr. Russo

> Mr. Russo takes 3 min. less than Mr. Lloyd to pack a case when each works alone. One day, after Mr. Russo spent 6 min. in packing a case, the boss called him away, and Mr. Lloyd finished packing in 4 more minutes. How many minutes would it take Mr. Russo alone to pack a case?

Saturated Fats

> One vegetable oil contains 6% saturated fats and a second contains 26% saturated fats. In making a salad dressing how many ounces of the second may be added to 10 oz. of the first if the percent of saturated fats is not to exceed 16%?

They will use a "work problem" schema for the first and a "mixture problem" schema for the second. More experienced math students would include both of these problems in a "linear equations" schema.

Several years ago, I did some studies of the imagery people use to solve elementary math problems (Hayes, 1973). When I gave people long-division problems to do in their heads, I heard my subjects do a lot of talking to themselves: "Two-seventy-three into nine-forty-one, is two, and two times two-seventy-three is . . ." "Aha!" I said to myself. "Auditory images are important here." What really surprised me though was the behavior of subjects recruited from the faculty of the modern languages department. These subjects were people who were born in Europe but had been in the United States for many years and spoke excellent English. These subjects did a lot of talking, too, but in French, Spanish, Italian, Polish, or Latvian—whatever language they spoke when they originally learned division. One person told me that he did elementary mathematics in Catalan, his first language, and more advanced mathematics in Spanish, the language he used about quite undeliberately, without any effort, and have no conscious purpose or use. For example, I don't use them deliberately to help me remember the songs. They rather just accompany my listening to a song, or singing it, or thinking about it, as a sort of image of the song. As the tune unfolds, I mentally proceed along the route."

We have observed that people use very diverse forms of representation even when they are solving simple problems. Numbers may be represented as the sound of words in one's native language. They may be represented as visual images

Table 1. One Subject's Color Associations

of print forms, Braille patterns, or colors. And they may be represented as finger movements. Most of our subjects used two or more of the forms. Thus, when they are representing a problem, people appear to have considerable choice in how they represent it. This choice is important because the form of representation that subjects choose can make a big difference in the difficulty they have in solving problems, and in the success they have in generalizing the solutions.

External Representations

In many cases, an external representation is very helpful for solving problems. Drawing a sketch, jotting down lists, writing out equations, and making diagrams can help us to remember information and to notice new relations in the problem.

Consider the following rate problem:

> A car can average 20 mph up to Pike s Peak and 60 mph back down the same road. What is the average speed for the whole trip?

Some people will find this problem easy enough to solve in their heads. Others feel much more comfortable with pencil and paper—writing down relations as they occur to them and *not* trying to juggle all the facts in their heads at once. The scratch sheet of such a person might look like this:

$$(1)\ \text{average rate} = \frac{\text{total distance}}{\text{total time}}$$

$$(2)\ \text{distance up} = \text{distance down} = X$$

$$(3)\ \text{total distance} = 2X$$

$$(4)\ \text{time} = \frac{\text{distance}}{\text{rate}}$$

$$(5)\ \text{time up} = \frac{X}{20\ \text{mph}};\ \text{time down} = \frac{X}{60\ \text{mph}}$$

$$(6)\ \text{total time} = \frac{X}{20} + \frac{X}{60}$$

$$(7)\ \text{average rate} = \frac{2X}{\frac{X}{20} + \frac{X}{60}} = \frac{2X}{\frac{60X + 20X}{20 \cdot 60}} = \frac{20 \cdot 60 \cdot 2X}{80X}$$

$$(8)\qquad = \frac{20 \cdot 60 \cdot 2}{80} = \frac{2400}{80} = 30\ \text{mph}$$

Clearly this external representation is an enormous aid to memory. The problem solver can compute total time in lines 4 through 6 without having to remember total distance. In computing average rate in lines 7 and 8, he can apply each algebraic step without having to remember the effects of previous steps. Working without such an external representation would be very difficult for most people.

Other kinds of external representation can also be very useful memory aids. For example, matrix representation is very useful in solving identification problems such as this one:

Dickens, Einstein, Freud, and Kant

Dickens, Einstein, Freud, and Kant are professors of English, Physics, Psychology, and Philosophy (though not necessarily respectively).

1. Dickens and Freud were in the audience when the psychologist delivered his first lecture.
2. Both Einstein and the philosopher were friends of the physicist.
3. The philosopher has attended lectures by both Kant and Dickens.
4. Dickens has never heard of Freud.

Match the protessors to their fields.

Our task is to match the professors to their fields. To do this, we construct a matrix as shown in Figure 5. Now, reading sentence 1, we conclude that the psychologist is neither Dickens nor Freud, so we put X's (indicating combinations ruled out) in two blocks as shown in the top matrix of Figure 6. In the second line, we learn that Einstein is neither the philosopher nor the physicist, and in the third line, that the philosopher is neither Kant nor Dickens, so we can fill in four more X's, as shown in the middle matrix of Figure 6. Now that leaves only Freud who could be the philosopher, so we put an O in the block corresponding to Freud and philosophy, and X out the other alternative fields for Freud (see remaining matrix, Figure 6). Proceeding in this way (though you may have some difficulty with the last few steps), you can identify the fields of all of the professors.

The matrix, like the notations in the previous problem, provides us with a great deal of help in remembering the results we have obtained in previous steps. Without such aids, some problems would be difficult or impossible to solve.

	English	Physics	Psychology	Philosophy
Dickens				
Einstein				
Freud				
Kant				

Figure 5. A matrix for the Dickens, Einstein, Freud, and Kant problem

TOTAL VOLUME AFTER = ALCOHOL AFTER + WATER AFTER

If any of these relations hadn't occurred to you after reading the problem text, the diagram could have given you a very useful hint.

Now, solve this problem before you proceed.

> A board was sawed into two pieces. One piece was two-thirds as long as the whole board and was exceeded in length by the second piece by 4 ft. How long was the board before it was cut?

Did you notice the contradictory nature of the problem? Paige and Simon found that people who draw a diagram to represent this problem can use the diagram to discover its contradictory nature. People who do not draw a diagram are likely to miss the contradiction and some may be quite happy to accept an answer of —12 feet for the length of the board!

External representations, then, can be enormously useful both in remembering the details of a problem and in understanding the relations among its parts. You should *always* consider using them when you are solving difficult problems.

Change and Growth in Presentations

An important fact about a representation is that it can change or develop as we work on the problem. Often enough, when we start to solve a problem, there are some important parts that we are vague about or which have escaped us entirely. We may not fully understand the whole problem until we have worked on it for some time. When people start to solve the Monster Problem, they usually have a pretty clear understanding of the initial state, the goal, and the operator. Often though, they don't really understand the restrictions. As they try to make a move, we may hear them mutter, "If two globes are holding the same monster . . . No. That's

This new, more precise representation of the goal can help you to avoid false leads in your search for a solution.

I observed another example of change in representation firsthand when a mend challenged me to solve the Four Knights Problem. This problem involves a 3 × 3 chess board and four chess pieces—two white knights and two black knights, arranged as shown in Figure 8. The goal is to interchange the positions of the white and black knights using only legal knight moves. For those who aren't familiar with chess, Figure 9 shows the legal knight moves. The knight can move one space straight ahead and one space diagonally forward.

Unfortunately, I had never seen the Four Knights Problem before. However, on general principles, I set up some guidelines in searching for a solution. First, I decided to work with an external representation of the problem to help me keep my place. I used a 3 × 3 matrix like that shown in Figure 8, on which I pencilled the current position of each piece and erased the previous position. Second, I knew that if I moved pieces at random, I would have trouble remembering which move. Now, because of the result I observed in the external representation, I added a larger operator to my representation—a macro-operator—consisting of eight knights' moves. Using the macro-operator, I could solve the problem in just two moves rather than 16.

Figure 10. Four Knights Problem Rotated One Step Clockwise From Original Position

Figure 11. Four Knights Problem Rotated Two Steps Clockwise From Original Position

Very often, then, we change our representation of a problem while we are solving it. In many cases, these changes appear to be improvements which make the problem easier to solve. If you are having difficulty in solving a problem, it makes sense to consider changing your problem representation. A useful way to proceed, as Polya (1945) has suggested, is to reexamine the problem statement very carefully. Perhaps we can make an inference which will help us represent the goal more accurately as in the Matchstick Problem. Or perhaps we can form a macro-operator as in the Four Knights Problem. Careful examination of each of the four problem parts—initial state, goal, operators, and restrictions-can suggest ways to improve our representation.

Representations Make a Difference

A problem may be difficult or impossible for us to solve in one representation, but much easier in another. For example, consider the Nine Dots Problem.

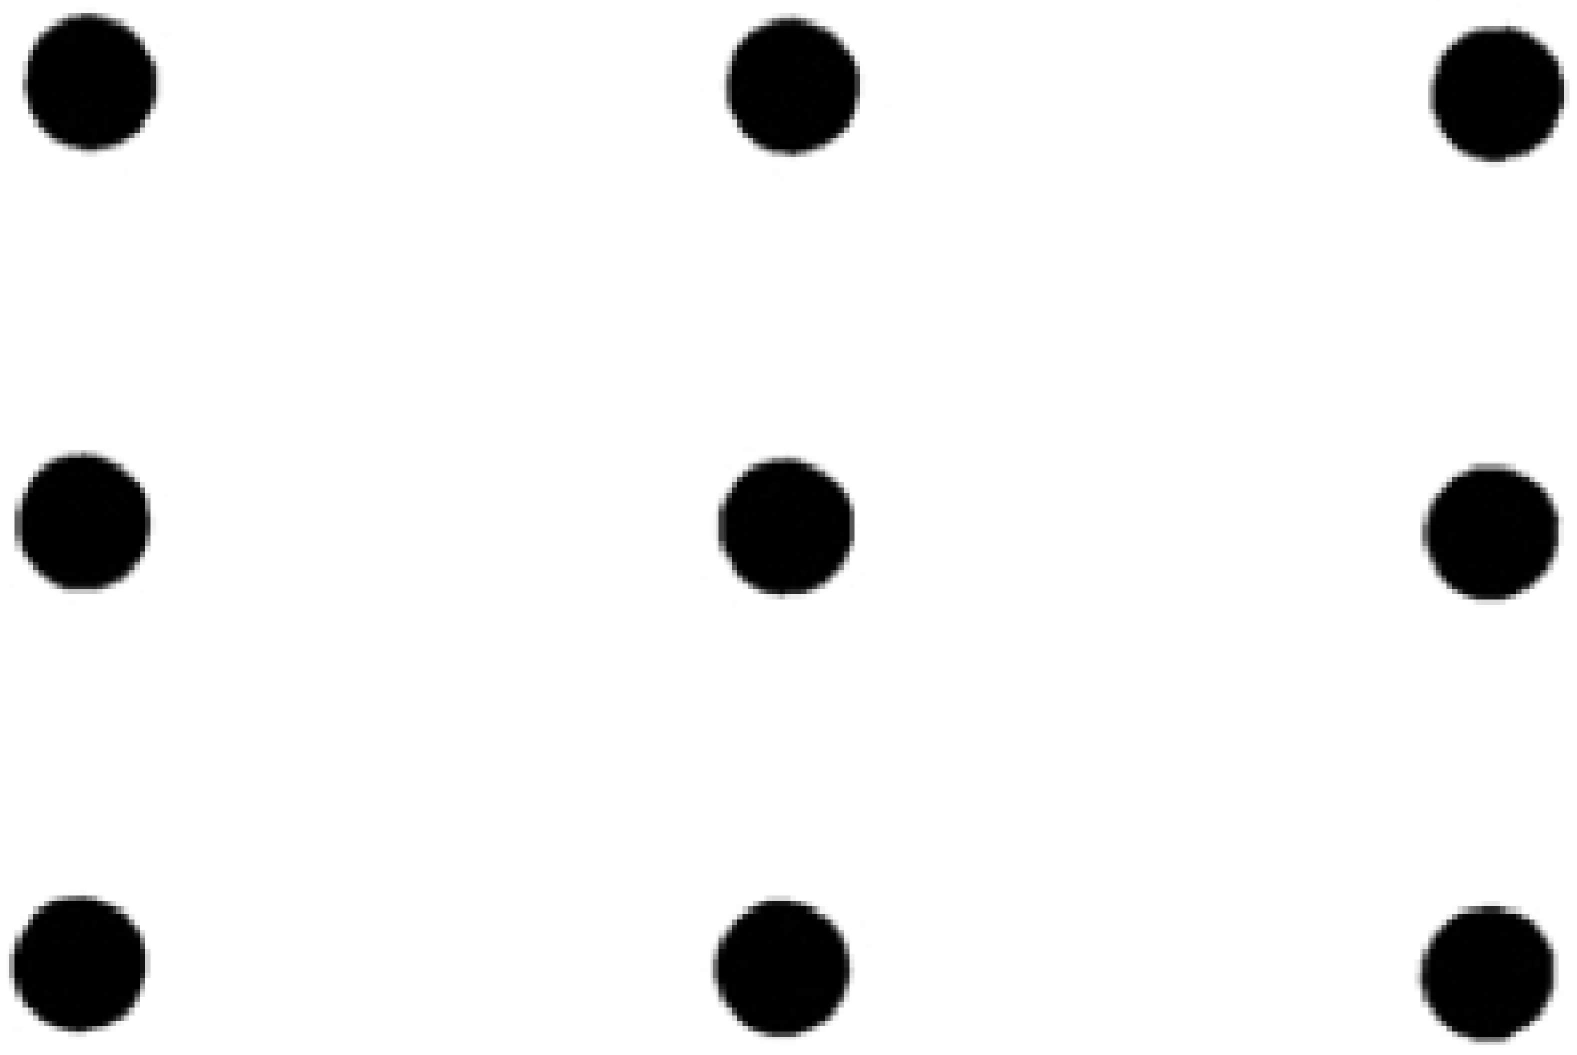

The Nine Dots Problem

> Without raising your pencil from the paper, draw four straight lines so that each of the dots above is touched by at least one of the lines.

If you don't already know the problem, try to solve it before proceeding. Some people have trouble because they have added a restriction to their representation which makes the problem unsolvable.* The restriction is that the lines should never extend beyond the square defined by the nine dots. Typical solution attempts for subjects adding this restriction are shown in Figure 12.

The first representation gave trouble because it added an extra restriction to the problem. Any representation that adds or deletes significant things from the initial state, from the goal, from the operators, or from the restrictions is very likely to give us serious trouble. We can avoid this trouble to some extent by checking our representation very carefully against the problem statement before we launch into any massive solution attempt.

Even when our representation of a problem is essentially correct, though, there are other ways to form a correct representation. Some of them are easier to use than others. Consider the following river-current problem.

A River-Current Problem

> You are standing by the side of a river which is flowing past you at the rate of 5 mph. You spot a raft 1 mi. upstream on which there are two boys helplessly adrift. Then you spot the boys' parents 1 mi, downstream paddling upstream to save them. You know that in still water the parents can paddle at the rate of 4 mph. How long will it be before the parents reach the boys?

One very natural way to represent the problem is to take the point of view of the observer standing by the side of the river. (The problem really sets you up to do this.) We can compute the speed of the boys with respect to the observer (5 mph downstream) and the speed of the parents with respect to the observer (5 mph—4 mph = 1 mph downstream). The difference in speed between the boys and their parents is four miles per hour. Thus, it should take half an hour to cover the two-mile distance.

An alternate and simpler way to represent the problem is to take the point of view of the boys on the raft. If we take this point of view, we can ignore the rate at which the boys and their parents are moving with respect to the observer. (The observer really *is* irrelevant in this problem.) In addition, we can ignore the rate of the current, since it is affecting the boys and their parents equally. (If this seems strange to you, remember that we routinely ignore

E