TEACHING PROBLEM-SOLVING SKILLS

ALAN H. SCHOENFELD

Department of Mathematics, Hamilton College, Clinton, NY 13323

1. Introduction and Overview. This paper deals with two questions:

A. Can we accurately describe the strategies used by "expert" mathematicians to solve problems? and

B. Can we teach students to use those strategies?

I make two basic assumptions. First: as a result of their problem-solving experience, mathematicians develop consistent and useful problem-solving strategies. Second: most students are not aware of, or do not use, these strategies. For example, consider the following problems.

Problem 1: Let a, b, c, and d be given numbers between 0 *and* 1. *Prove that*

$$(1-a)(1-b)(1-c)(1-d) > 1-a-b-c-d.$$

Problem 2: Determine the sum $\frac{1}{2!}+\frac{2}{3!}+\cdots+\frac{n}{(n+1)!}$.

Problem 3: Prove that if $2^n - 1$ *is a prime, then n is a prime.*

Ostensibly, all three of these problems are accessible to high school students. None of them require mathematical knowledge beyond algebra, and all of them have straightforward solutions. Yet college students and professional mathematicians attack these problems in dramatically different ways.

On Problem 1 most students will laboriously multiply the four factors on the left, subtract the terms on the right, and then try to prove that $(ab+ac+ad+bc+bd+cd-abc-abd-acd-bcd+abcd)>0$—usually without success. Virtually all of the mathematicians I've watched solving it begin by proving the inequality $(1-a)(1-b)>1-a-b$. Then they multiply this inequality, in turn, by $(1-c)$ and $(1-d)$ to prove the three- and four-variable versions of it.

Likewise in Problem 2, most students begin by doing the addition and placing all the terms over a common denominator. A typical expert, on the other hand, begins with the observation: "That looks messy. Let me calculate a few cases." The inductive pattern is clear and easy to prove.

The colleague who read Problem 3 and said, "That's got to be done by contradiction," was typical; given the structure of the problem, one really has no alternative. Yet this almost automatic expert observation is alien to students: a large number of those to whom I have given the problems either respond with comments like "I have no idea where to begin" or try a few calculations to see whether the result is plausible and then reach a dead end.

Of course, these are special problems for which expert and novice performance is remarkably consistent. While the experts did not consciously follow any strategies, their behavior was at least consistent with these "heuristic" suggestions:

a. For complex problems with many variables, consider solving an analogous problem with fewer variables. Then try to exploit either the method or the result of that solution.

b. Given a problem with an integer parameter n, calculate special cases for small n and look for a pattern.

Alan H. Schoenfeld received his Ph.D. from Stanford in 1973, working under Karel deLeeuw in topology and measure theory. After two years at Davis he joined the SESAME group at U.C. Berkeley, where he began his current "research and development in 'human AI,' where the research goal is to create detailed models of human problem-solving, the development goal to use the models to teach students to solve problems the way experts do." He would like to hear from colleagues who are either studying or trying to teach problem solving.

The author dedicates this paper to the memory of Karel deLeeuw, in appreciation of his spirit of inquiry and his devotion to his students.

c. Consider argument by contradiction, especially when extra "artillery" for solving the problem is gained by negating the desired conclusion.

Many of the novices were unaware of the strategies, and many others "knew of them" (that is, upon seeing the solution they acknowledged having seen similar solutions) but hadn't thought to use them. Expert and novice problem-solving are clearly different. The critical question is: Can we train novices to solve the problems as experts do?

My answer is a provisional "yes." I think it is possible to give a course in which we can teach students to solve a wide variety of problems—including problems unlike any solved in the course—better and more efficiently than they could otherwise. But there are many questions to be answered. How much sophistication and background do students need before such instruction can be effective? What does it take to understand a strategy like "establish subgoals" and how to use it? What do you need in addition to the mastery of individual strategies? Briefly, my thesis is this.

First, the strategies are more complex than their simple descriptions would seem to indicate. If we want students to use them, we must describe them in detail and teach them with the same seriousness that we would teach any other mathematics. Second, there is clear evidence that the strategies do make a difference—when there are only a small number of them and they are taught under closely controlled circumstances. Third, being able to use individual strategies is not enough: you have to know which ones to use, and when. We can provide students with a reasonable structure for efficient problem-solving and can actually demonstrate improvement.

2. The Complexity of Heuristic Strategies. The first person to describe problem-solving strategies in such a way that they could be *taught* (although he does not claim that they can be and makes no promises about the results) was Pólya. In *How to Solve It* (1945) and the two volumes of *Mathematical Discovery* (1962 and 1965) Pólya laid the foundation for explorations in heuristics.

Let us define a *heuristic strategy* as a general suggestion or technique which helps problem-solvers to understand or to solve a problem. Heuristic strategies include the "fewer variables," "calculating special cases," and "argument by contradiction" strategies described in section 1. Fig. 3 gives many more. Many investigators have attempted to show that these strategies can help students to solve problems. However, the results are generally inconclusive, in part because these apparently simple strategies can turn out to be very complex. Consider the following strategy and a few problems.

"To solve a complicated problem, it often helps to examine and solve a simpler analogous problem. Then exploit your solution."

Problem 4: Two points on the surface of the unit sphere (in 3-space) are connected by an arc A which passes through the interior of the sphere. Prove that if the length of A is less than 2, then there is a hemisphere H which does not intersect A.

Problem 5: Let a, b, and c be positive real numbers. Show that not all three of the terms $a(1-b)$, $b(1-c)$, and $c(1-a)$ can exceed $1/4$.

Problem 6: Find the volume of the unit sphere in 4-space.

Problem 7: Prove that if $a^2+b^2+c^2+d^2=ab+bc+cd+da$, then $a=b=c=d$.

These four problems, like Problem 1, can be solved by the "analogous problem" strategy. Yet it is unlikely that a student untrained in using the strategy would be able to apply it successfully to many of these. Part of the reason is that the strategy needs to be used *differently* in the solution of each problem.

In solving Problem 1, we built up an inductive solution from the two-variable case, using the result of the analogous problem as a stepping stone in the solution of the original.

In contrast, analogy is used in Problem 4 to furnish the idea for an argument. The problem is hard to visualize in 3-space but easy to see in the plane: we want to construct a diameter of a unit circle which does not intersect an arc of length 2 whose endpoints are on the circle. Observing that the diameter parallel to the straight line between the endpoints has this property enables us to return to 3-space and to construct the analogous plane.

Problem 5 is curious. It looks as though the two-variable analogy should be useful, but I haven't found an easy way to solve it. At first the one-variable version looks irrelevant, but it's not. If you solve it, and think to take the product of the three given terms, you can solve the given problem. So again we exploit a result, but this time a different result in a different way.

Problem 6 exploits both the methods and results of the lower-dimensional problems. We integrate cross-sections using the same method; the measures of the cross-sections are the results we exploit.

In Problem 7 it would seem apparent that the two-variable problem is the appropriate one to consider. However, *which* two-variable problem is not at all clear to students. A large number of those I have watched tried to solve:

Problem 7′: Prove that $a^2+b^2=ab$ *implies that* $a=b$, *instead of*

Problem 7″: Prove that $a^2+b^2=ab+ba$ *implies* $a=b$.

We conclude that the description "exploiting simpler analogous problems" is really a convenient *label* for a collection of similar, but not identical, strategies. To solve a problem using this strategy, one must (a) think to use the strategy (this is nontrivial!), (b) be able to generate analogous problems which are appropriate to look at, (c) select among the analogies the appropriate one, (d) solve the analogous problem, and (e) be able to exploit either the method or the result of the analogous problem appropriately.

This strategy isn't especially complex. "Look for an inductive solution when you see an integer parameter" is easier, but "establish and exploit subgoals" is far more difficult. The moral of this section is that it's easy to underestimate the amount of work that would go into teaching students to use even a single strategy. We should single it out as a useful strategy; we should give sample problems (like the ones above) showing how it works; we should remind the students of its use in other problems when we use it; and we should chide them when they fail to use it.

3. Teaching Strategies to Students Makes a Difference (Sometimes). There is some convincing evidence that "experts" do use the kind of heuristic strategies we have been discussing. But these strategies are rarely taught explicitly: they might be called "good habits learned through experience in problem solving." Thus some might claim (indeed, have claimed) that any gain in a problem-solving environment is not due to the strategies taught but to the practice in solving problems. For that reason I conducted an experiment in which two groups of students received essentially the same problem-solving training, except that the strategies were mentioned explicitly to only one of the groups.

Each of the seven upper-division science majors I worked with was taught and tested individually. "Instruction" was provided on tape recordings, so that it was replicable and could be checked by colleagues. The students were trained to solve problems "out loud" and the tests were tape-recorded. The experiment was designed to test the students' use of five particular heuristic strategies:

1. Draw a diagram if at all possible.
2. If there is an an integer parameter, look for an inductive argument.
3. Consider a logical alternative: arguing by contradiction or contrapositive.
4. Consider a similar problem with fewer variables.
5. Try to establish subgoals.

During the instruction, each of the students worked twenty problems and then saw solutions to each of them. Each had the same amount of time for problem solving and for seeing the solutions. However, four of the students received a heuristic "extra." They were given a list and explanation of the five strategies used in the experiment and an "overlay" to each solution explaining how the strategy had been used. Fig. 1 gives the solution to a problem we have already considered. The right-hand side is the solution seen by all students. The left-hand side was seen only by the "heuristics" students. All the students were reminded periodically (during practice sessions and tests) to review carefully what they were doing, and the reminders to the heuristics group included the phrase "Look over the list of strategies."

The four-variable problem is too complicated.

You are given the real numbers a, b, c, and d, each of which lies between 0 and 1. Prove the inequality

$$(1-a)(1-b)(1-c)(1-d)>1-a-b-c-d.$$

Solution

Can we learn something from a similar one-variable problem? No.

Suppose we start by proving the equation

How about the comparable two-variable problem? It's easy to solve.

$$(1-a)(1-b)>1-a-b. \tag{*}$$

If we multiply out the left, (*) is true if and only if

$$1-a-b+ab>1-a-b,$$

which is true if and only if $ab>0$. But $ab>0$, since we were given that a and b are both positive. This proves (*). Now let's build on this. The number c is between 0 and 1, so $(1-c)$ is positive. Multiplying both sides of (*) by $(1-c)$, we get

Can we use the result? Yes . . . build up to 3 variables . . . Then build up to 4.

$$(1-a)(1-b)(1-c)>(1-a-b)(1-c), \text{ or}$$

$$(1-a)(1-b)(1-c)>1-a-b-c+ac+bc.$$

Since ac and bc are both positive, we obtain

$$(1-a)(1-b)(1-c)>1-a-b-c. \tag{**}$$

Remember, when a problem is complicated...consider a similar problem with fewer variables. Then try to use either the method or the result to solve the original problem.

Continuing in the same vein, we notice that $(1-d)$ is positive; multiplying both sides of (**) by $(1-d)$, we obtain

$$(1-a)(1-b)(1-c)(1-d)>(1-a-b-c)(1-d), \text{ or}$$
$$(1-a)(1-b)(1-c)(1-d)>1-a-b-c-d+ad+bd+cd.$$

As before, we see that ad, bd, and cd are all positive. Thus

$$(1-a)(1-b)(1-c)(1-d)>1-a-b-c-d,$$

which is what we wanted to prove.

FIG. 1

I give a detailed description of the results in "Explicit heuristic training..." Even this small sample yielded a statistically significant difference in pretest to post-test gains: each of the four "heuristics" students outscored the non-heuristics students, a one-in-35 chance. But, more important, the transcripts of the solutions show that explicit use of the strategies accounted for the differences between the two groups.

The "fewer variables" practice included Problems 1, 5, and 7, as given above, and the following.

Problem 8: Show that it is impossible to find real numbers a,b,c,d,e,A,B,C,D,E such that

$$x^2+y^2+z^2+r^2+s^2=(ax+by+cz+dr+es)(Ax+By+Cz+Dr+Es)$$

for all values of x, y, z, r, and s.

For each of these, all of the students saw how the one- or two-variable analog was used to solve the original problem. If "practice" is what counts, all of them should have solved this post-test problem:

Problem 9: Suppose p, q, r, and s are positive real numbers. Prove the inequality

$$\frac{(p^2+1)(q^2+1)(r^2+1)(s^2+1)}{pqrs} \geqslant 16.$$

All four of the "heuristics" students solved it, but only one of the others did. The other two non-heuristics students multiplied through by $pqrs$ and tried (unsuccessfully) to deal with the resulting inequality.

This shows that we cannot rely on students' abilities to grasp useful problem-solving strategies when the students are not given explicit instructions on their use and that the instruction "made a difference." More precisely, we can say that heuristics made a difference under experimental conditions in which (1) there were a limited number of strategies to "worry" about, (2) there were periodic reminders to consider using the strategies, and (3) the test problems were clearly amenable to one of the suggested "heuristic" approaches. Taking heuristics instruction out of the laboratory and into the real world will be no easy task.

4. The Need for Global Strategies. If heuristics can make a difference, why are the results in the literature so equivocal? The major studies (see those by Goldberg, Kantowski, Lucas, Smith, Webb, Wilson) are generally encouraging, but that's about all. I can see two possible reasons for this, one of them easily remedied.

We have seen that it is easy to underestimate the amount of work required to teach a particular strategy. For example, Kantowski reported at the 1978 NCTM meetings that students in a problem-solving experiment failed to "look back" over their solutions, in spite of the fact that 40 per cent of instruction time was spent "looking back." Videotapes of the class sessions showed that after each problem the teacher had stepped aside and said, "Now let's look at what we've done," and proceeded to do so. But the value of the strategy was not stressed. Students were not shown why it is useful to "look back"—so they didn't. If we want students to take a strategy seriously, we have to convince them of its usefulness.

But even if we succeed in teaching students to use a series of important heuristic strategies, I see no guarantee that there will be clear signs of improvement in their general problem solving. Knowing how to use a strategy isn't enough: the student must think to use it when it's appropriate. To justify this claim I argue first by analogy, then briefly describe a supportive experiment.

We can think of a heuristic strategy as a "key" to unlock a problem. There are a large number of such "keys," and a given problem may be "openable" by only a few of them. Imagine facing a locked door with a key ring on which there are thirty keys, two of which will open the door. If you only have time to try three or four keys in the lock, the fact that the "right" key is somewhere on the chain may not help you very much. On the other hand, a strategy for selecting the right key might. If you could narrow down the collection of "candidate" keys to ten, the opportunity to try three or four of these gives you a much better chance of success.

Consider techniques of integration in elementary calculus. There are fewer than a dozen important techniques, all of them algorithmic and relatively easy to learn. Most students can learn integration by parts, substitution, and partial fractions as individual techniques and use them reasonably well, as long as they know which techniques they are supposed to use. (Imagine a test on which the appropriate technique is suggested for each problem. The students would probably do very well.) When they have to select their own techniques, however, things often go awry. For example,

$$\int \frac{x\,dx}{x^2-9},$$

a "gift" first problem on a test, caused numerous students trouble when they tried to solve it by partial fractions or, even worse, by a trigonometric substitution!

In "Presenting a Strategy for Indefinite Integration" (this MONTHLY, Oct. 1978) I discuss an experiment in which half the students in a calculus class (not mine) were given a strategy for selecting techniques of integration, based on a model of "expert" performance. The other students were told to study as usual—using the miscellaneous exercises in the text to develop their own approaches to problem solving. Average study time for members of the "strategy" group was 7.1 hours, while for the others it was 8.8 hours; yet the "strategy" group significantly out-performed the rest on a test of integration skills—in spite of the fact that they were not given training in integration, just in selecting the techniques of integration.

The "moral" to the experiment is that students who cannot choose the "right" approach to a problem—even in an area where there are only a few useful straightforward techniques—do not perform nearly as well as they "should." If we leap from techniques of integration to general mathematical problem-solving, the number of potentially useful techniques increases substantially, as does the difficulty and subtlety in applying the techniques. An efficient means for selecting approaches to problems, for avoiding "blind alleys," and for allocating problem-solving resources in general thus becomes much more critical. Without it, the benefits of training in individual heuristics may be lost.

5. The Model. The model of "expert" performance described below serves as the framework for my courses in problem solving. Of course, any attempt to characterize mathematical problem solving on just a few sheets of paper must leave out much more than it includes.

The global outline of the strategy is given in Fig. 2. I use a flow chart to indicate the generally dynamic but structured nature of the process; it is meant as a guide to profitable behaviors, not as a straitjacket that orders and restricts them. Various individual heuristics often come into play most appropriately at certain phases of the process. These are listed in Fig. 3.

The problem-solving process begins with an *analysis* of what the problem entails. This includes getting a "feel" for the problem by looking for what is given, what is asked for, why the "givens" are given, whether what is asked for seems plausible, what major mechanisms seem to apply, what mathematical context the problem fits into, and so on. Which heuristics (if any) are brought to bear during analysis may depend on both the problem and who is solving it (how much of a "problem" or routine exercise is this "task" to the individual?). But examples of the appropriate use of some heuristic strategies at this stage of problem-solving are:

(1) to draw a diagram even when the problem appears amenable to a different kind of argument, such as in the following:

Problem 10: Find those values of t for which the equations

$$\left\{\begin{array}{c} x^2-y^2=0 \\ \textit{and} \\ (x-t)^2+y^2=1 \end{array}\right\} \textit{have } 0, 1, 2, 3, \textit{or } 4 \textit{ solutions.}$$

(2) to examine special cases and try to solve them or to determine patterns. For example, in:

Problem 11: Given $a, b>0$, *determine* $\lim_{n\to\infty}(a^n+b^n)^{1/n}$.

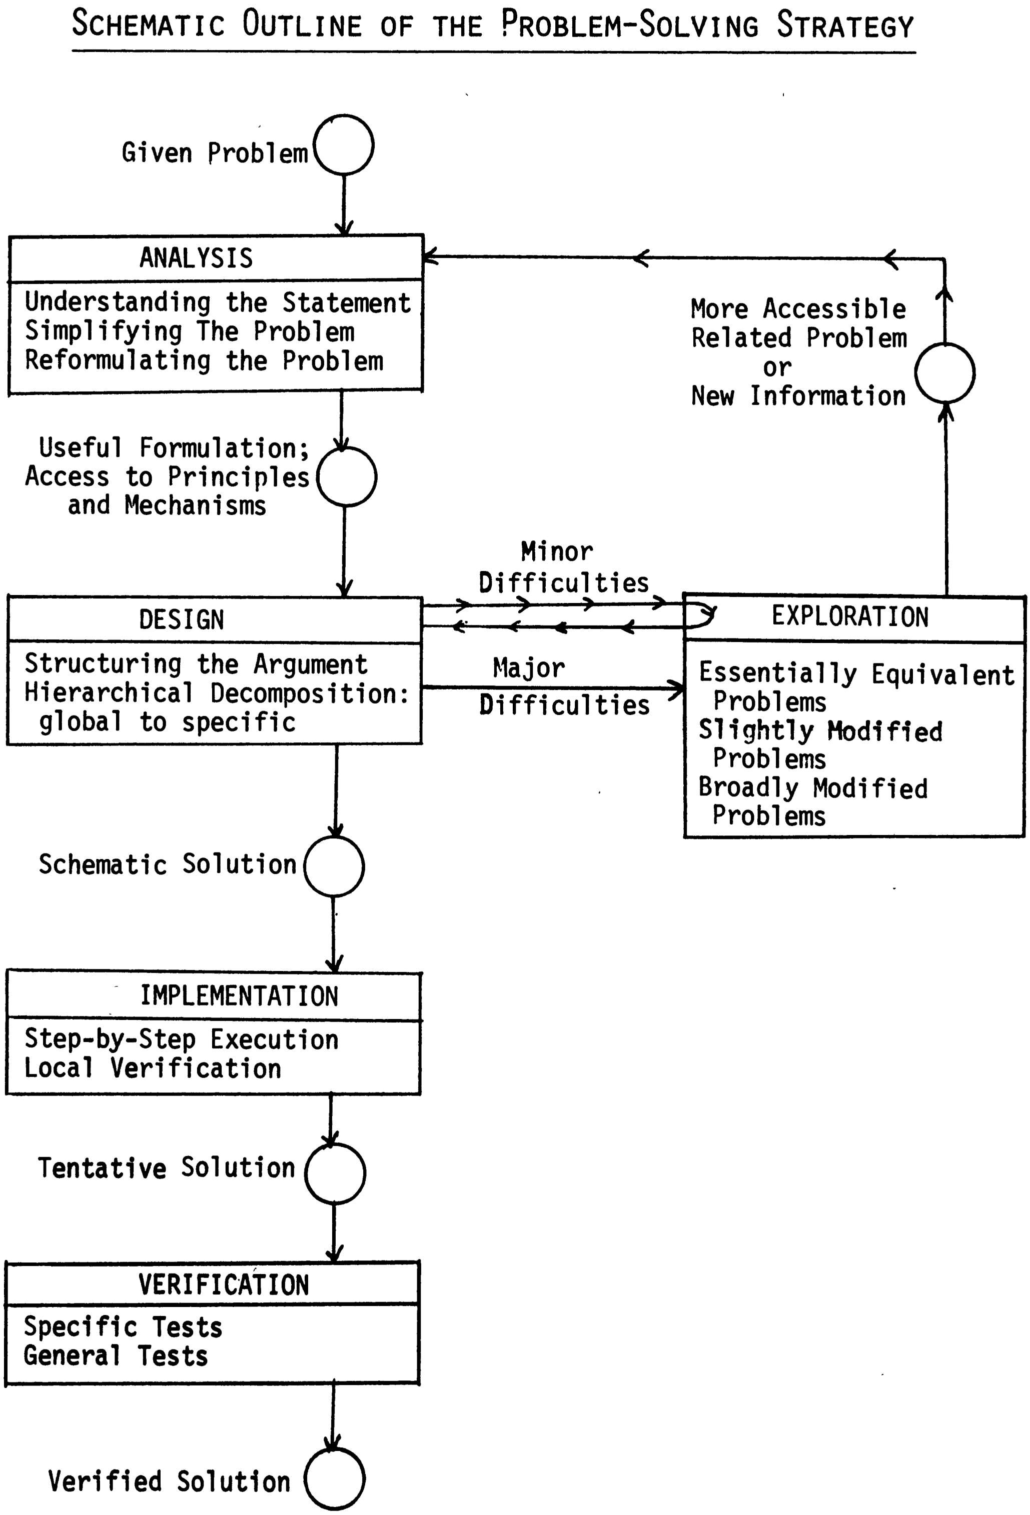
SCHEMATIC OUTLINE OF THE PROBLEM-SOLVING STRATEGY
Given Problem
ANALYSIS
Understanding the Statement
Simplifying The Problem
Reformulating the Problem
More Accessible Related Problem or New Information
Useful Formulation; Access to Principles and Mechanisms
Minor Difficulties
DESIGN
Structuring the Argument
Hierarchical Decomposition: global to specific
Major Difficulties
EXPLORATION
Essentially Equivalent Problems
Slightly Modified Problems
Broadly Modified Problems
Schematic Solution
IMPLEMENTATION
Step-by-Step Execution
Local Verification
Tentative Solution
VERIFICATION
Specific Tests
General Tests
Verified Solution

FIG. 2

Frequently Used Heuristics

Analysis

1. Draw a diagram if at all possible.
2. Examine special cases:
   a. Choose special values to exemplify the problem and get a "feel" for it.
   b. Examine limiting cases to explore the range of possibilities.
   c. Set any integer parameters equal to $1,2,3,\ldots,$ in sequence, and look for an inductive pattern.
3. Try to simplify the problem by
   a. exploiting symmetry, or
   b. "without loss of generality" arguments (including scaling).

Exploration

1. Consider essentially equivalent problems:
   a. Replacing conditions by equivalent ones.
   b. Re-combining the elements of the problem in different ways.
   c. Introduce auxiliary elements.
   d. Re-formulate the problem by
      (i) change of perspective or notation
      (ii) considering argument by contradiction or contrapositive
      (iii) assuming you have a solution and determining its properties.
2. Consider slightly modified problems:
   a. Choose subgoals (obtain partial fulfillment of the conditions)
   b. Relax a condition and they try to re-impose it
   c. Decompose the domain of the problem and work on it case by case.
3. Consider broadly modified problems:
   a. Construct an analogous problem with fewer variables.
   b. Hold all but one variable fixed to determine that variable's impact.
   c. Try to exploit any related problems that have similar
      (i) form
      (ii) "givens"
      (iii) conclusions.

   Remember: when dealing with easier related problems, you should try to exploit both the *result* and the *method of solution* on the given problem.

Verifying Your Solution

1. Does your solution pass these specific tests?
   a. Does it use all the pertinent data?
   b. Does it conform to reasonable estimates or predictions?
   c. Does it withstand tests of symmetry, dimension analysis, and scaling?
2. Does it pass these general tests?
   a. Can it be obtained differently?
   b. Can it be substantiated by special cases?
   c. Can it be reduced to known results?
   d. Can it be used to generate something you know?

FIG. 3

one might want to set $a=1$; in:

Problem 12: Find $\sum_{n=1}^{\infty} 1/n(n+1)$.

one might want to compute the sums for 1, 2, 3, 4, and 5 terms to see the (surprisingly obvious) answer.

(3) to look for preliminary simplifications. In:

Problem 13: Find the largest area of any triangle which can be inscribed in a circle of radius R.

one might (i) consider first the unit circle, (ii) note that, without loss of generality, one can assume that the base of the triangle is horizontal, and (iii) examine several sketches and try to guess an answer before jumping into an analytic solution.

Design is in a sense a "master control." It is not really a separate box on the flow chart but rather pervades the entire solution process. Its function is to ensure that the problem solver is engaged in activities most likely to be profitable. It entails keeping a global perspective on the problem and proceeding hierarchically. An outline of the solution should be developed at a rough qualitative level and then elaborated in detail as the solution process proceeds. For example, detailed calculations or complex operations should not be performed until (i) alternatives have been explored, (ii) there is clear justification for them, and (iii) other stages of the problem solution have proceeded to the point where the results of the calculations either are necessary or will clearly prove useful. (How painful it is to expend time and energy solving a differential equation only to discover that the solution is of no real help in the "next" global phase of the problem!)

Exploration is the heuristic "heart" of the strategy; it is in the exploratory phase that most of the problem-solving heuristics come into play. Fig. 3 shows that exploration is divided into three stages. Generally, the suggestions in the first stage are either easier to employ or more likely to provide direct access to a solution of the original problem than those in the second stage; likewise for the relation between stages 2 and 3. All other factors being equal, the problem solver in the exploration phase would briefly consider those suggestions in stage 1 for plausibility, select one or more, and try to exploit it. If the strategies in stage 1 prove insufficient, one proceeds to stage 2; if need be, when stage 2 has been exhausted one tries the strategies in stage 3. If substantial progress is made at any stage, the problem solver may either return to *design* a plan for the balance of the solution, or may decide to re-enter *analysis*, with the belief that the insights gained in *exploration* can help re-cast the problem in a way not previously seen.

Implementation needs little comment, save that it should be the last step in the actual problem solution. *Verification*, on the other hand, deserves attention if only because it is so often slighted. At a local level, one can catch silly mistakes. At a global level, a review of the solution can yield alternative methods, show connections to other seemingly unrelated subject matter, and, on occasion, clarify a useful technique that then can be incorporated into one's global problem-solving approach.

6. The Instruction and Some Results. The model described in the previous section has served as the foundation for two courses in problem solving—one given to eight upper-division mathematics majors at Berkeley in 1976, and one given to nineteen (mostly) lower-division liberal arts students at Hamilton College in 1979. In each course the students were given the model as a guide to the problem-solving process. Each class session was devoted to a series of problems solvable by one (or more) of the strategies listed in Fig. 3. We would go over the solution, stressing both the use of the particular strategy and (relative to the model) how to approach the whole problem with some efficiency.

There were substantial differences between the two courses, largely because of the difference in mathematical sophistication between the two groups of students. For example, the suggestion to consider the argument "by contradiction or contrapositive" meant very different things to the two groups. For the upper-division mathematics majors, a few examples and a discussion of when it might be appropriate to consider the strategy were enough; I was essentially "pulling together" in coherent form what they had seen used as a tool in a variety of places. It was an entirely different story with the freshmen. Many of them were unconvinced that there is a need to prove things mathematically *at all*, and many had never seen an argument by contradiction! I had assigned the following as part of a take-home midterm at Berkeley.

Problem 14: Let a be a digit from 1 to 9. Which numbers of the form aaaa...a, where a is repeated n times, are perfect squares?

All but one student solved it. In contrast, it occupied us (in bits and pieces) in class for a number of days at Hamilton.

There were, however, many similarities between the two groups. Few of the students had had conscious access to any of the heuristics strategies we have been discussing. During the first class session at Berkeley, only two of the eight students succeeded in finding $\Sigma_{n=1}^{\infty} 1/n(n+1)$, using the telescoping series. None had thought to try values of 1,2,3,4 for n. Similarly, they did not draw diagrams where appropriate, etc. At the end of the course there was clear evidence that the students were consciously using heuristics effectively and recognizing the appropriateness of particular heuristics to particular types of problems. For example, this was on the final examination.

Problem 15: Let S be any nonempty finite set. We define E(S) to be the number of subsets of S which have an even *number of elements, including the null set and possibly S. Determine E(S) in closed form for any finite set S, and prove your answer.*

Seven of the eight students approached the problem by looking for an inductive pattern, a far cry from their entering behavior. (The eighth student, the only one who claimed to have seen the problem before, outlined a combinatoric argument.)

There were similar results for this strategy in the course at Hamilton. On a test at the beginning of the course, 4 students of 19 thought to calculate sums in this:

Problem 16: What is the sum of the first 89 *odd numbers?*

and some others used Gauss's pairing of terms (which they had seen before) to get the answer. On the final examination, 18 of 19 solved the following.

Problem 17: What is the sum of the coefficients of $(x+1)^{31}$?

Nothing resembling Problem 17 had been discussed in the course.

Likewise in both courses there were clear differences on the "fewer variables" strategy and on other strategies that are equally well defined. Student performance on Problem 1, which I used at the beginning of both courses, has been discussed. Problem 9 was on both final exams, and more than three-fourths of the students in each course solved it.

I should balance these "success stories" somewhat. As we saw in section 2, heuristics are subtle and students can easily go astray when trying to use them. There we saw that choosing the "right" analogous problem was not easy. Also, we "experts" have the ability to "see through" certain forms which even the more advanced undergraduates are unable to recognize. The last line in our proof of Problem 7 at Berkeley was: "Since $(a-b)^2+(b-c)^2+(c-d)^2+(d-a)^2=0$, then $a=b$, $b=c$, $c=d$, and $d=a$." Leaving this line on the board, I gave the class the following problem.

Problem 18: Let the numbers a_i and b_i be given for $i=1,2,\ldots,n$. Determine necessary and sufficient conditions on the a_i and b_i such that there exist real numbers A and B such that $(a_1x+b_1)^2+\cdots+(a_nx+b_n)^2=(Ax+B)^2$ for all x.

The morass of symbols in the second problem, including the variable x, the subscripts, and so on, obscured the similarity between the two problems. The students failed to see the essentially analogous structure that "a sum of squares equals something that is or can be made equal to zero." They were thus unable to solve the second problem.

We cannot expect students to use any heuristic in ways that go significantly beyond the way they have been shown to use it. Asked to find the number of positive integer divisors $D(N)$ of the integer N, my students had no trouble in seeing that they should calculate $D(N)$ for different values of N. When the results did not look suggestive, however, they did not think to ask, "What values of N give particular values of $D(N)$?"—which unlocks the problem. Likewise problems that call for a clever synthesis of two heuristics they have studied (like Pick's theorem, which calls for first fixing one variable and then doing an induction on the free one) will often prove beyond the students' reach. These are not grounds for despair, but merely a call for realistic expectations.

Our reasonable expectations can actually be rather high. At least in the short term, testing before and after the course indicates some substantial progress on the part of the students. Of course, the more important question is the long-term impact of the instruction and the effect, if any, that it has on the students' performance outside the class. It's still too early to tell, but preliminary reports from students who have taken the courses have been enthusiastic and favorable.

To be perfectly honest, I should mention that a course in problem solving requires a substantial commitment from all concerned. The teacher has to be especially flexible, because it's the *process* of problem solving that counts and the teacher is essentially serving as a "coach" to the students. The students are being asked to *think*, and to create, rather than to "recite" subject matter. That's not an easy task, but it is a critically important one—and ultimately a very rewarding one, well worth the effort on the part of the students. It is also, of course, a source of tremendous gratification for the successful instructor.

Note. While this article was in press, I taught the problem-solving course again—this time with an extensive barrage of before-and-after testing and a "control group" for comparison. The results were quite dramatic and will be written up in later reports.

I thank Jean J. Pedersen for valuable editorial criticism.

Bibliography

A. References

B. S. Bloom, Problem Solving Processes of College Students, University of Chicago Press, Chicago, Ill., 1950.

G. D'Amour and C. Wales, Improving problem-solving skills through a course in guided design, Engineering Education, 67 (5) (1977).

P. Daniels, Strategies to Facilitate Problem Solving, Cooperative Research Project No. 1810, Brigham Young University, Provo, Utah, 1964.

J. Dodson, Characteristics of Successful Insightful Problem Solvers, No. 71-13, 048, University Microfilms, Ann Arbor, 1970.

K. Duncker, On Problem Solving, Psychological Monographs, 58, no. 5, American Psychological Association, Washington, D.C., 1945.

E. Flaherty, Cognitive Processes Used in Solving Mathematical Problems, No. 73-23, 562, University Microfilms, Ann Arbor, 1973.

D. Goldberg, The Effects of Training in Heuristic Methods on the Ability to Write Proofs in Number Theory, No. 75-07836, University Microfilms, Ann Arbor, 1975.

J. W. Hadamard, Essay on the Psychology of Invention in the Mathematical Field, Dover, New York, 1954.

J. R. Hayes, Memory, goals, and problem solving, in B. Kleinmuntz, ed., Problem Solving: Research, Method, and Theory, Wiley, New York, 1966.

J. Kilpatrick, Problem solving in mathematics, Review of Educational Research, 1969.

______, Research in the Teaching and Learning of Mathematics, paper presented to the MAA, Dallas, 1973.

______, Variables and Methodologies in Research in Problem Solving, paper for research workshop on Problem Solving in Mathematics Education, University of Georgia, 1975.

I. Lakatos, Proofs and Refutations, Cambridge University Press, Cambridge, 1976; rev. ed., 1977.

L. N. Landa, Algorithmization, Educational Technology Publications, Englewood Cliffs, N.J., 1974.

______, The ability to think: How can it be taught? Soviet Education, vol. 18, 5, March 1976.

J. Larkin, Skilled problem solving in physics: A hierarchical planning model, Journal of Structural Learning (in press).

J. Lucas, An Exploratory Study on the Diagnostic Teaching of Heuristic Problem-Solving Strategies in Calculus, No. 72-15, 368, University Microfilms, Ann Arbor, 1972.

A. Newell and H. Simon, Human Problem Solving, Prentice-Hall, Englewood Cliffs, N.J., 1972.

G. Pólya, How to Solve It, Princeton University Press, Princeton, N.J., 1945.

______, Mathematical Discovery, vols. 1 and 2, Wiley, New York, 1962 (vol. 1) and 1965 (vol. 2).

A. Schoenfeld, Presenting a strategy for indefinite integration, this MONTHLY, 85 (1978) 673–678.

______, Can heuristics be taught?, in J. Lockhead, ed., Cognitive Process Instruction, Franklin Institute Press, Philadelphia, 1979.

______, Explicit heuristic training as a variable in problem solving performance, Journal for Research in Mathematics Education, May 1979.

______, Heuristics in the Classroom, 1980 Yearbook of the National Council of Teachers of Mathematics, 1980.

B. F. Skinner, Teaching Thinking, Chapter 6 of The Technology of Teaching, Appleton-Century-Crofts, New York, 1968.

P. Smith, The Effect of General vs. Specific Heuristics in Mathematical Problem Solving Tasks, No. 73-26367, University Microfilms, Ann Arbor, 1973.

R. Turner, Design problems in research on teaching strategies in mathematics, in T. Cooney, ed., Teaching Strategies: Papers from a Research Workshop, ERIC, Columbus, Ohio, 1976.

W. Wickelgren, How to Solve Problems, W. H. Freeman, San Francisco, 1974.

D. R. Woods et al., How can one teach problem solving?, Program for Instructional Development Newsletter, Ontario Universities, May 1977.

B. Problem Sources

J. C. Burkill and H. M. Cundy, Mathematical Scholarship Problems, Cambridge University Press, Cambridge, 1961.

S. J. Bryant et al., Non-Routine Problems in Algebra, Geometry, and Trigonometry, McGraw-Hill, New York, 1965.

E. B. Dynkin et al., Mathematical Problems: An Anthology, Gordon and Breach, New York, 1969.

H. Eves and E. P. Starke, The Otto Dunkel Memorial Problem Book, Mathematical Association of America, 1957.

D. Lidsky et al., Problems in Elementary Mathematics (transl. by V. Volosov), Mir Publishers, Moscow, 1963.

G. Pólya and J. Kilpatrick, The Stanford Mathematics Problem Book, Teacher's College Press, New York, 1974.

E. Rapaport, transl., The Hungarian Problem Book, Mathematical Association of America, 1963.

D. O. Shklarsky et al., The USSR Olympiad Problem Book, W. H. Freeman, San Francisco, 1962.

C. W. Trigg, Mathematical Quickies, McGraw-Hill, New York, 1967.

Yaglom and Yaglom, Challenging Mathematical Problems with Elementary Solutions, Holden-Day, San Francisco, 1967.

THE FORMULA OF FAA DI BRUNO

STEVEN ROMAN

Department of Mathematics, University of South Florida, Tampa, FL 33620

1. Introduction. Almost every calculus student is familiar with the formula of Leibniz for the nth derivative of the product of two functions

$$D^n f(t)g(t) = \sum_{k=0}^{n} \binom{n}{k} D^k f(t) D^{n-k} g(t).$$

A much less well known formula is that of Faà di Bruno for the nth derivative of the composition $f(g(t))$ (see Theorem 2). It is the purpose of this paper to give a new proof of this formula.

Several proofs of this formula have appeared in the literature. For example, in [1] there is a

The author received his Ph.D in mathematics in 1975 at the University of Washington under Branko Grünbaum. Since then he has been Instructor of Applied Mathematics at the Massachusetts Institute of Technology, Lecturer of Mathematics at the University of California at Santa Barbara, and Assistant Professor of Mathematics at the University of South Florida. He is currently Associate Professor of Mathematics at California State University at Fullerton. His research interests include combinatorics, graph theory, and applied mathematics. —*Editors*