

Problem solving: the difference between what we do and what we tell students to do*

George M. Bodner

Purdue University, West Lafayette, IN 47907

e-mail: gmbodner@purdue.edu

Introduction

It is slightly over 30 years since I was first asked to teach something known as ‘general chemistry’ at the University of Illinois. Without any idea of what went into that course, the order in which topics should be taught, or the amount of time that should be devoted to each topic, I asked a couple of senior colleagues what they did when they taught this course and tried to do the same.

During the course of that first semester, I found that I enjoyed teaching and that the students enjoyed having me as their instructor. Everything was going well until I made the mistake of analyzing the students’ answers to the exams I gave them. I was shocked; or, in the language of Jean Piaget, utterly disequibrated. In spite of clear, concise, well-organized, and well-delivered lectures, I found that bright, hardworking science and engineering majors couldn’t solve ‘simple’ problems on topics that had been taught — and taught well!¹ Thus, it shouldn’t be surprising that one of the topics I became interested in as a beginning researcher in chemical education was problem solving.

Over the course of about 20 years, the author has worked with roughly a dozen graduate students pursuing M.S. or Ph.D. degrees in chemical education whose studies focused on different aspects of problem solving. It is the results of these students’ work that serves as the basis for this paper.

Problem-solving research


Virtually all of our insight into problem solving has come from research that uses qualitative methods, in

which we interview people struggling to solve problems and ask them to ‘think aloud’ — to talk about what they are doing, or what they are thinking, while they are involved in the problem-solving process. We’ve worked in a variety of courses, from general chemistry through the sophomore organic and inorganic courses, to physical chemistry, and even advanced organic chemistry courses taken by graduate students.

A few samples of the kinds of questions we have given to participants in our interviews are shown below. The first question is from an early study of problem solving by science and engineering majors enrolled in a general chemistry course at Purdue; the second is from a study of undergraduates, graduate students, and faculty trying to predict the product of an inorganic reaction; the third comes from a study of students enrolled in a graduate-level organic chemistry course; and the fourth is from a study of graduate students’ understanding of aspects of 2D FT NMR.

- Uranium reacts with fluorine to produce a compound, which is a gas at 57°C. The density of this gas is 13.0 g/L at 57°C and 1 atm pressure. Is the molecular formula of this compound (a) UF₂, (b) UF₃, (c) UF₄, (d) UF₅ or (e) UF₆?
- Predict the products of the following reactions:
$$\begin{array}{ll} \text{Na} + \text{H}_2\text{O} \rightarrow & \text{NaOH} + \text{Cl}_2 \rightarrow \\ \text{MgO} + \text{H}_2\text{O} \rightarrow & \text{H}_2\text{S} + \text{Cl}_2 \rightarrow \\ \text{Ba}_3\text{N}_2 + \text{H}_2\text{O} \rightarrow & \text{NaOH} + \text{SO}_2 \rightarrow \\ \text{XeF}_2 + \text{D}_2\text{O} \rightarrow & \text{NO}_2 + \text{H}_2\text{O} \rightarrow \end{array}$$
- Explain the following reaction (Figure 1).

* This paper is based on the Royal Society of Chemistry’s 2003 Nyholm Lecture given by the author.



- Starting from thermal equilibrium (with M^0 aligned along the Z axis) and assuming no delays between pulses, predict in which plane the magnetization vector M will lay after experiencing the following pulse sequence. Assume the RF transmitter is aligned along the +X axis. $90^\circ_x, 90^\circ_x, 180^\circ_{x'}, 90^\circ_{x'}, 270^\circ_{x'}, 90^\circ_{x'}, 90^\circ_x$.

engine versus the caboose — such as the wavy line indicating smoke escaping from the engine of the trains in the following drawing. So far, without exception, they have all labeled the length of one train as ‘x’ and the other as ‘3x’. They have also labeled the distance between the engine of one train and the engine of the other, and between the caboose on one train and the caboose on the other, as shown in Figure 2.


Problems versus exercises

Chemists, who are used to differentiating between metals and nonmetals, between ionic and covalent bonds, between acids and bases, between polar and non-polar solvents, and so on, should be particularly sensitive to the role that duality can play in describing a phenomenon. Thus, they shouldn’t be surprised that early research on problem solving was driven, in part, by attempts to distinguish between the way subject-matter experts and novices approached certain tasks.² Our work has led us to question the value of comparing the work of experts and novices because we don’t believe a given task means the same thing to both groups.³ To illustrate this, consider a problem we have given to hundreds, if not quite thousands, of industrial chemists participating in workshops on problem solving or graduate students participating in a training program for teaching assistants.

Two trains are stopped on adjacent tracks. The engine of one train is 1000 yards ahead of the engine of the other. The end of the caboose of the first train is 400 yards ahead of the end of the caboose of the other. The first train is three times as long as the second. How long are the trains?

Let’s assume, for the sake of convenience, that the two trains are headed in the same direction. Let’s also remind ourselves of the definition of a *caboose* — the car that used to be placed at the end of a train, which was used by the crew on the train.

We’ve found that industrial chemists invariably start with a drawing, using some convention to identify the



They then write an equation in one unknown and solve for ‘x.’

$$3x + 400 = x + 1000$$

$$2x = 600$$

$$x = 300$$

The teaching assistants do virtually the same things. The only fundamental difference between the two groups is the tendency for those in industry to write ‘x = 300’ and for those in academia to write ‘x = 300 yd.’

When the industrial chemists are told that there is no partial credit in this course, and they therefore get a zero, they get mad. They get a zero for the obvious reason — they haven’t answered the question! When told that they are going to receive no credit for their answer, the graduate students shrug this off. They’re


U.Chem.Ed., 2003, 7, 38

use to not getting the credit they feel they deserve on exams.


For now, let's focus on two observations about this problem. First, when faced with a novel problem, practicing chemists almost always start with a drawing of some kind, and frequently annotate the drawing with relevant information. Second, practicing chemists stop their problem solving activities when they get to the point that they fully understand the problem; not when they get the 'answer'.

Now let's consider another question:

What is the molarity of an acetic acid solution, if 34.57 mL of this solution is needed to neutralize 25.19 mL of 0.1025 M sodium hydroxide.⁴



What would you expect practicing chemists to do? Would they start with an equation or formula, such as: $n = M \times V$? or with a drawing such as Figure 3?



The answer should be obvious — in the absence of explicit instruction to do so, no practicing chemist would draw a picture when doing this routine exercise. They would all start by feeding numbers into an equation.

These examples suggest that a given individual might exhibit fundamentally different behaviors on different problem-solving tasks. To help the reader understand the source of these differences, we need to define the terms *problem* and *problem solving*. We'll start with John Hayes' definition of the term *problem*.⁵

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.


According to Hayes, the presence of a gap between where you are and where you want to be is a necessary — but not sufficient — criterion to classify a task as problem. There also has to be an element of uncertainty or confusion, if not downright ignorance, about how one is going to cross that gap.

Almost 20 years ago, Wheatley proposed a definition of *problem solving* that is consistent with Hayes' definition of a problem. Wheatley argued that problem solving is "*what you do, when you don't know what to do*".⁶

If the definitions proposed by Hayes and Wheatley are accepted, it should be easy to understand why we stress the difference between tasks that are *routine exercises* and those that are *novel problems*. When people first encounter these terms, they often assume that the difference between an exercise and a problem is based on difficulty, or complexity. Our work has shown that problems are neither inherently more difficult nor more complex. The only difference between an exercise and a problem is the element of familiarity.

Consider the following question from a general chemistry exam.


What weight of oxygen is required to burn 10.0 grams of magnesium?



This is a routine exercise for a practicing chemist, but a novel problem for students who encounter chemistry for the first time.

Another example of this phenomenon can be found in the following question from a sophomore organic chemistry course.

Robinson annulation reactions involve two steps: Michael addition and aldol condensation. Assume that Michael addition leads to the following intermediate.



U.Chem.Ed., 2003, 7, 39

What would be produced when this intermediate undergoes aldol condensation?

This is a problem for most chemists, but a routine exercise for those who either teach or do organic chemistry. Not because they're any brighter, but because they are so familiar with similar tasks.

The distinction between exercises and problems is important because it is a potential source of miscommunication between instructors and their students. We tend to put a content expert in the classroom for whom tasks that arise during the semester are routine exercises, and expect that individual to 'teach' students for whom these tasks are novel problems. Consider what would happen if we asked organic chemists to teach physical chemistry, or vice versa. The approach they would take to 'teaching' students how to solve problems would be different, not simply because of differences in the way they think about chemistry, but because of differences in their familiarity with these tasks.

The difference between the way exercises and problems are worked is particularly well illustrated by the examples that appear in so many textbooks. These examples have several characteristics.

- They are logical sequences of steps.
- They string together in a linear fashion.
- They proceed from the initial information to the solution.

These textbook solutions, which are often mirrored by instructors in the classroom, are examples of a phenomenon that can be called 'forward-chaining' or 'forward-working'. They are examples of how routine exercises are worked by individuals with many years of experience with similar tasks. However, they have little, if any, similarity to the approach successful problem solvers use when they encounter novel problems.

As Herron once noted:⁷

"The solutions given by authors in textbooks bear little resemblance to what experts do when they work unfamiliar problems. (Textbook solutions ... describe the most efficient pathway to a solution and probably represent how an expert who solves such problems routinely would approach the task.)"

Herron and coworkers have argued that:⁸

"... [textbook] examples must convey to the

students an unrealistic idea about how problems are actually attacked. The examples provide no indication of the false starts, dead ends, and illogical attempts that characterize problem solving in its early stages, nor do they reveal the substantial time and effort expended to construct a useful representation of a problem before the systematic solution shown in examples is possible."

Instead of comparing the work of experts working on routine exercises with novices struggling with novel problems, we have chosen a different duality. We prefer contrasting the work of successful problem solvers (of any age) with the behavior of those who are less successful when these individuals encounter problems that are outside of their area of expertise.

Models of problem solving

One of the goals of our work is the development of a model of problem solving that has two characteristics. First, and foremost, it must fit our experimental data from interviews with successful problem solvers working on what is, for them, a novel problem. Second, it must be 'teachable'; it must be a model that can be given to students that can improve their problem solving performance in chemistry.

Let's therefore look at several models of problem solving that have been proposed, starting with Polya's model that consists of four stages.⁹

- Understand the problem
- Devise a plan
- Carry out the plan
- Look back

This model makes sense. It seems logical that we would start by understanding the problem, then devising a plan, then carrying out the plan, and then looking back to check our work and consolidate our gains.

Unfortunately, Polya's model is not consistent with our work. To try to convince the reader of this, consider the following problem which is based on the experimental data collected when one of the first xenon fluoride compounds was analyzed.¹⁰

A sample of a compound of xenon and fluorine was confined in a bulb with a pressure of 24 torr. Hydrogen was added to the bulb until the pressure was 96 torr. Passage of an electric spark through the mixture produced Xe and HF. After

the HF was removed by reaction with solid KOH, the final pressure of xenon and unreacted hydrogen in the bulb was 48 torr. What is the empirical formula of the xenon fluoride in the original sample?

We've given this problem to practicing chemists who do not teach general chemistry. They inevitably get the correct answer, but analysis of what they say when we ask them to work this problem out loud suggests that they do not follow the four stages of Polya's model. Indeed, a common comment heard when they finally get to the answer — XeF_4 — is: "Oh, it's an empirical formula problem!" In other words, our experience with this problem — like the 'two trains' problem cited earlier — suggests that the process of problem solving is over when one gets to the point that they understand the problem.

Several other models of problem solving that are logical extensions of Polya's model have been discussed elsewhere.¹¹ They all have the disadvantage of not being consistent with the patterns we've observed for successful problem solvers working on novel problems. Let's therefore turn to a model proposed by Alex Johnstone and co-workers.¹² This model assumes that each learner has a working-memory capacity (X) and that each problem has a working-memory demand (Z), which is defined as the maximum number of steps activated by the least able individual.

The Johnstone-El Banna model assumes that when the working-memory capacity of the individual is equal to or larger than the demand on working memory ($X \geq Z$), we have a necessary, but not sufficient, condition for success. It isn't sufficient because success also depends on prior knowledge; on whether the prior knowledge is easily accessible; on the student's motivation (inclination, interest, etc.); and so on.

This model assumes that students won't be successful when the demand on working memory exceeds the capacity of working memory ($Z > X$), unless the student can organize the demand on working memory so that it is smaller than his or her working-memory capacity. Johnstone and co-workers note that when the demand of the problem exceeds capacity, there is a sharp drop in performance. But, some students ($\approx 10\%$) seem to be able to solve problems for which the demand exceeds capacity ($Z > X$) because of chunking devices that reduce the demand on working

memory.

Let's assume, for the moment, that the Johnstone-El-Banna model is correct when it is applied to situations that meet the six criteria proposed by Tsaparlis.¹³ Furthermore, let's assume that Niaz is correct when he concludes that: "Teachers can facilitate success by decreasing the amount of information required for processing, and thereby avoiding working memory overload".¹⁴ Now what? From the perspective of this model, there isn't much we can do to improve student performance in our classes — other than helping them learn how to 'chunk' information. We simply have to accept the limitations our students bring to the classroom, and conclude that the only way we can improve their performance is to lower the intellectual rigor of the tasks we give them.

We believe that we can do more than this. Based on research on problem solving in mathematics, Wheatley proposed an anarchistic model of problem solving that describes what successful problem solvers do when they work on novel problems.⁶ As noted most recently by Calimsiz,¹⁵ this model is consistent with the results of our problem-solving interviews.

An Anarchistic Model of Problem Solving

- Read the problem
- Now read the problem again
- Write down what you hope is the relevant information
- Draw a picture, make a list, or write an equation or formula to help you begin to understand the problem
- Try something
- Try something else
- See where this gets you
- Read the problem again
- Try something else
- See where this gets you
- Test intermediate results to see whether you are making any progress toward an answer
- Read the problem again
- When appropriate, strike your forehead and say, "son of a ..."
- Write down 'an' answer (not necessarily 'the answer')
- Test the answer to see if it makes sense
- Start over if you have to, celebrate if you don't

U.Chem.Ed., 2003, 7, 41

“Draw a Picture”

There are several stages in this model that deserve explicit attention. In the ‘two trains’ problem, we saw the role that a drawing that is annotated with relevant information can play in solving a novel problem. We’ve also seen, in the calculation of the molarity of the acetic acid solution, that drawings aren’t done when people encounter a routine exercise.

Over the years, several of the author’s colleagues have noted how difficult it is to get their students to “draw something” while working on problems in organic chemistry. We’ve encountered a similar resistance among juniors taking physical chemistry, often because they can’t visualize the system they are working with.

In a study of problem solving by graduate students and early career faculty within the context of combined spectra interpretation, Cartrette¹⁶ noted that successful problem solvers in this study were much more likely to draw out molecular fragments as they were deduced in the problem solving process — in other words, to “draw something.”

Our experience suggests that one cannot get students to draw a picture as a routine part of their problem solving process by telling them that *they should do* this. We’ve found that students are more receptive to including this step when we tell them that this is something that *we do*.

“Try Something”

We’ve often described the steps “try something” and “try something else” as ‘playing with the problem’. Unfortunately, far too many of our students — particularly those who are struggling with the course — believe you can’t ‘play’ with a problem. They believe that ‘trial and error’ is not a legitimate strategy for problem solving — often because they haven’t seen any of THEIR instructors use this strategy in class. This is unfortunate because trial and error seems to be one of the most powerful strategies that our successful problem solvers own.

There is abundant evidence in our data that successful problem solvers routinely encounter a cue during problem solving that causes them to ask: Am I getting anywhere? Many beginners forget to do this. They exhibit a ‘garden-path syndrome’, working the problem the way they might walk through a garden — smelling the roses along the way, but not noticing that

they aren’t getting anywhere. Successful problem solvers tend to start over when they find that they aren’t making any progress toward the answer; beginners often fail to do this.

“Does the Answer Make Sense?”

The penultimate step in this model is particularly important. We have found that beginners seldom test their answers to see if they make sense for several reasons. First, they’ve never seen anyone do this when they’ve watched their instructors work out the solutions to tasks that are exercises for the instructors. Second, they are seldom given the information they would need to do this.

Whenever we think about the penultimate stage in this model we are reminded of the phenomenon known as a Fermi calculation or Fermi estimate. Enrico Fermi had a reputation for asking students at the University of Chicago questions that seemed impossible and then showing them how to use common knowledge to estimate the answer. (For our purposes, “to test the answer to see if it makes sense.”)

The most oft-cited example of a Fermi calculation involves asking students to estimate the number of piano tuners in Chicago. One starts with an estimate of the population of Chicago, the fraction of this population who are likely to own pianos, the frequency with which pianos are tuned, and so on.

Fermi calculations can be relatively simple, such as estimating how long it would take to eat your weight in food (about a month), or how much trash produced in a typical house each year (about 1000 pounds). But they can also be considerably more challenging, such as estimating the fraction of the continental U.S. covered by automobiles.

The Difference Between Exercises and Problems

In summary, we would like to argue that Polya’s model is an ideal approach to working a routine exercise. One reads the question, understands the task, devises a plan, and so on. We might go so far as to argue that one of the characteristic tests of whether a task is an exercise is to ask: How is the solution found? Exercises are worked in a linear, forward-chaining, rational manner. Our model suggests that problem solving is cyclic, reflective, and can appear irrational. Experts who watch students struggle with a problem are tempted to intervene; to show the

U.Chem.Ed., 2003, 7, 42

'correct' way of obtaining the answer. This makes the expert feel good, but it doesn't necessarily help the individual who is struggling with the problem for the first time.

Our experience suggests that the anarchistic model can be taught to students and that students who pick up this approach to problem solving often do better in the course than those who do not. Teaching this model involves presenting it to the students at the beginning of the semester and then explicitly using it, over and over again, throughout the course.


We've found that this model is equally productive in introductory courses taken by freshman from the various schools that require a year of general chemistry and in physical chemistry courses taken by junior chemistry and biology majors. In a recent study of sophomores taking a year-long course in organic chemistry, Calimsiz¹⁵ found that this model most closely reflected the process by which successful students worked problems that asked them to propose a set of reactions that would transform a given starting material into a given product.

Implications for teaching

For some time, we've been recommending that instructors who teach introductory chemistry courses draw a picture for every task they work in class from the beginning of the Fall semester until the end of the Spring semester. We find that when this is done, the number of 'C'-s in the class goes down and the number of 'B'-s and 'A'-s goes up. An example of this phenomenon is shown below:


Question: Use the K_{sp} for calcium fluoride to calculate its solubility in grams per liter.

Figure 4 contains the information extracted from the statement of the problem, as well as information that



is derived while one proceeds toward the answer. Part of its power comes from the fact that it is a symbolic representation. Part comes from the fact that it is a second representation, which often brings to the students' attention details that are not always as obvious as we might expect. (Such as the fact that there are twice as many F^- ions as Ca^{2+} ions in the solution.) It is also important to recognize that this diagram is a chunking device, as called for in the Johnstone-El-Banna model. It brings together information, thereby reducing the demand on working memory.

In a prior publication in this journal¹⁷ we looked at the implications of this idea when it is applied to the kind of descriptive chemistry one finds in modules on inorganic chemistry. The example we used in that paper was based on years of watching what happens when TA's try to explain the reaction between the triiodide ion and thiosulfate.



Transfer of learning

Gage and Berliner argue that "*the transfer of skills, knowledge, learning strategies, etc., is a fundamental goal of all levels of education*".¹⁸ Transfer has been defined as the "*use of information or skills characteristic of one domain or context in some new domain or context*".¹⁹ Transfer can occur from one problem to another within a course; from one class to another; from one year to another; from school to home; and from school to work.²⁰

For years, one of the Author's goals has been building problem solving skills that transfer to other courses and, eventually, to improvements in on-the-job performance. It is for that reason that he includes

questions such as the following in his textbooks because he believes that these questions provide the basis for practising the anarchistic model of problem solving he incorporates into his classes.


In 1773 Benjamin Franklin observed that one teaspoon of oil spilled on a pond near London spread out to form a film that covered an area of about 22,000 square feet. If a teaspoon of oil has a volume of about 5 cm^3 and the oil spread out to form a film roughly 1 molecule tall, what is the average height of an oil molecule?⁴

It is also the reason that he collects questions such as the following, which was proposed for a course on environmental problem solving.

In a remote area in Nepal, the concentration of aluminium in outdoor air at ground level averages $9.4 \times 10^{-8} \text{ } \mu\text{g}/\text{cm}^3$. (It is much higher inside the Sherpa dwellings because of wood and yak dung burning). At the same site, the Al concentration in the top 1 cm of fresh snow averages $0.12 \text{ } \mu\text{g}/\text{g}$, while in the top 1 cm of three-day-old snow it averages $0.20 \text{ } \mu\text{g}/\text{g}$. Calculate the average deposition velocity of the Al falling to the ground when it is not snowing.²¹

For many years, a family of problems has circulated that were the primary concern of individuals doing research on problem solving and the limited number of people who collected puzzles for sheer intellectual excitement. Consider the following problem, for example (Figure 5).

Each of the following cards has a letter on one side and a number on the other. Which card or cards would you have to turn over in order to find out whether the following rule is true or false: If a card has a vowel on one side, it has an even number on the other side.²²



It is possible that someone might get the answer to this question without creating one of the drawings discussed in this paper, but it is not likely. Most people assume that the card with the 'A' on it must be turned over, and they are correct. What is surprising

is the relatively small number of people who assume that one also has to turn over the card with the '7' on it, to make sure that there is no vowel on the other side. The author is convinced that people who get this wrong would do better if they didn't try to solve the problem in their heads. If they were forced to keep records of their thought process while they thought about each card, one at a time; if they were forced to draw a picture or make a list to help them understand the problem.

This question surfaced recently because it is being used by companies such as Microsoft as part of the process by which potential employees are screened.²³ Other questions that are asked during these interviews include:

There are three ants at the three corners of a regular triangle. Each ant starts moving in a straight line toward another, randomly chosen corner. What is the probability that none of the ants collide?

Once again, it might be possible to get the right answer (one in four) without a drawing, but most people would have to start by translating the problem into a drawing.

One of the more popular questions in the Microsoft collection is the following:

Suppose you have eight billiard balls (or jars of pills, etc.). One of them is defective — it weighs more than the others. How do you tell, using a two-pan balance, which ball is defective in two weighings?

Many people to whom we've given this question conclude that either it can't be done or at least they can't do it. Everyone we've watched get the right answer has taken the approach of 'playing with the problem'. If you start by trying to put four balls on each pan of the balance you find that it can't be done in two weighings. So try putting three balls on each pan. If the pans balance, the defective ball is among the two you didn't weigh, and you can determine which one it is in a single weighing. If the pans don't balance, select the three balls that are heavier. Now try something else. Pick two of these three balls and put one of them on each pan. If one is heavier, it is the defective one. If the pans are in balance, then the ball that wasn't chosen must be defective.

The anarchistic model of problem solving presented in this paper was based on interview data on

U.Chem.Ed., 2003, **7**, 44

mathematically oriented problems. As we have continued our work, we've found that it applies just as well to non-mathematical problems such as organic synthesis or spectral interpretation. More recently, we've found that it applies to problems that extend beyond the domain of chemistry, and might therefore involve skills that are worth building and transferring.

References

1. G.M. Bodner, *J. Chem. Ed.*, 1986, **63**, 873.
2. J. Larkin, J. McDermott, D.P. Simon, and H.A. Simon, *Science*, 1980, **208**, 1335.
3. G.M. Bodner, A View From Chemistry, in *Toward A Unified Theory of Problem Solving: Views From the Content Domains*, M.U. Smith, Ed., Lawrence Erlbaum Associates: Hillsdale, NJ, **1991**, pp. 21-34.
4. G.M. Bodner and H.L. Pardue, *Chemistry: An Experimental Science*, 2nd Ed., John Wiley: New York, 1994.
5. J. Hayes, *The complete problem solver*. The Franklin Institute: Philadelphia, 1980.
6. G.H. Wheatley, Problem solving in school mathematics. MEPS Technical Report 84.01, School Mathematics and Science Center, Purdue University, West Lafayette, IN, 1984.
7. J.D. Herron, *The Chemistry Classroom: Formulas for Successful Teaching*, American Chemical Society: Washington, DC, 1996.
8. J.D. Herron and T.J. Greenbowe *J. Chem. Ed.*, 1986, **63**, 526.
9. G. Polya, *How to solve it: A new aspect of mathematical method*. Princeton University Press: Princeton, NJ, 1945.
10. H.F. Holtzlaw, W.R. Robinson, and W.H. Nebergall, *General Chemistry*, 7th Ed. D.C. Heath: Lexington, MA, 1984.
11. G.M. Bodner and J.D. Herron, *Problem Solving in Chemistry*, in *Chemical Education: Research-Based Practice*, J.K. Gilbert, Ed., Kluwer Academic Publishers, 2002.
12. Johnstone, A.H.; El-Banna, H. *Educ. Chem.*, 1986, **23**, 80.
13. G. Tsaparlis *Intern. J. Sci. Ed.*, 1998, **20**, 335.
14. M. Niaz *J. Chem. Ed.*, 1987, **64**, 502.
15. S. Calimsiz *How Undergraduates Solve Organic Synthesis Problems: A Problem Solving Model Approach*, Unpublished M.S. Thesis, Purdue University, 2003.
16. D. Cartrette *Using Combined Spectral Analysis to Probe the Continuum of Problem Solving Ability*, Unpublished Ph.D. Dissertation, Purdue University, 2003.
17. G.M. Bodner and D.S. Domin, *U. Chem. Ed.*, 2000, **4**, 22.
18. N.L. Gage and D.C. Berliner, *Educational Psychology*, 6th Ed. Boston, Houghtlin Mifflin, 1998.
19. A.F. Johnson and G.M. Bodner, Abstracts of the 226th ACS National Meeting, New York, NY, 2003.
20. A. Robins, *Connection Science*, 1996, **8**, 185.
21. J. Harte, *Consider a Spherical Cow: A Course in Environmental Problem Solving*, University Science Books: Sausalito, CA, 1988.
22. P. Johnson-Laird and P. Wason, in P. Johnson-Laird and P. Wason (Eds), *Thinking: Readings in Cognitive Science*, p. 143-157. Cambridge: Cambridge University Press, 1977.
23. W. Poundstone, *How Would You Move Mount Fuji?* Little, Brown and Company: New York, 2003.