Published in *Educacion Quimica*, **2005**, *16*(2), 222-229

A CULTURAL APPROACH TO PROBLEM SOLVING

George M. Bodner and Gautam Bhattacharyya
Purdue University, West Lafayette, IN 47907

ABSTRACT

For more than 20 years, our research group has been studying why bright, hard-working students often struggle to solve problems they encounter in undergraduate chemistry courses. Our work has spanned the breadth of the sub-disciplines of chemistry, from organic chemistry to physical chemistry. This paper will begin with a review of some of the general conclusions of this work that revolve around the importance of recognizing the difference between routine exercises and novel problems and then examine how the consequences of this difference should inform teaching and the evaluation of the various models of problem solving that have been proposed. We will conclude this review by introducing the reader to a new path our work has taken which suggests that science should be viewed as a culture, not merely a set of concepts and principles held together by a content-specific language. The implications of this path for both teaching and research will also be discussed.

INTRODUCTION

Introductory organic chemistry exams often ask students to write the mechanism for common reactions, such as the acetal formation reaction shown in the first part of Figure 1. When grading these exams, instructors often become dismayed that, in spite of showing comparable — if not identical — examples in class and discussing the mechanism of this reaction shown in the second part of Figure 1, the students often cannot get even the first step of the mechanism correct, the protonation of the ketone. The third part of Figure 1 depicts a common error that students make when they try to write the mechanism for this reaction.

Insert Figure 1 here

Analogous scenarios are encountered so often by instructors across the spectrum of chemistry courses that it isn't surprising that these individuals express their frustrations by questioning the amount of time students spend studying for exams, their work ethic, and/or their problem-solving skills. Our research group has spent more than twenty years trying to understand why even "good" students, who exhibit all the requisite study habits and skills, appear unable to apply what they have learned in class to problems they encounter on exams, in the laboratory, or as practicing chemists. These research efforts have spanned the breadth of the sub-disciplines of chemistry, from general chemistry through the sophomore organic and inorganic courses, to physical chemistry, and even graduate-level organic chemistry courses. This paper will start with an overview of some of the general conclusions that can be extracted from our work that has been described, in part, in a paper based on the 2003 RSC Nyholm lectures [1] and conclude with a brief discussion of a new path that our work has taken.

Most 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 talk about what they are doing or what they are thinking while they are involved in this process. Another useful source of information has been the analysis of answers to exam questions coupled with informal discussions with students about why they gave a particular answer when they took the exam.

PROBLEMS VERSUS EXERCISES

To help the reader understand one of the insights we've developed from our work, consider a problem that has been given to perhaps as many as a thousand people during workshops on problem solving for practicing chemists or during 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?

Regardless of whether they are practicing chemists working in industry or teaching assistants at the beginning of their graduate work, we have found that virtually everyone to whom we give this problem does essentially the same thing. They start with a drawing, in which they use some convention to identify the engine versus the caboose. They typically label the length of one train as "x" and the other as "3x." They label the distance between the engine of one train and the engine of the other; between the caboose on one train and the caboose on the other as shown in Figure 2.

Insert Figure 2 here

They then write an equation in one unknown, solve for "x", and report the answer.

$$3x + 400 = x + 1000$$
$$2x = 600$$
$$x = 300$$

The only fundamental difference between the two groups is the tendency for those in industry to write "x = 300" and for those in academics to write "x = 300 yd."

For now, we would like to 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.

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 consider a second 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?
>
> $$CH_3CO_2H(aq) + NaOH(aq) \rightarrow Na^+(aq) + CH_3CO_2^-(aq) + H_2O(l)$$

Would practicing chemists start with a drawing such as Figure 3, or would they start by writing an equation or formula, such as: $n = M \times V$?

Insert Figure 3 here

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

These examples suggest that a given individual may exhibit fundamentally different behaviors when presented with different problem-solving tasks. To understand these differences we need to define the terms *problem* and *problem solving*. The following is a traditional definition of the term "problem" proposed by John Hayes 25 years ago [2].

> 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 existence of a gap, by itself, does not guarantee that a task is a "problem;" there also must be an element of uncertainty about the method that can or will be used to bridge that gap. The two key elements in Hayes' definition of a problem were captured in Wheatley's definition of *problem solving* as "what you do when you don't know what to do" [3]. If one accepts the validity of these definitions one is lead to the logical conclusion that there is a fundamental difference between the way an individual responds to a *routine exercise* that is similar, but not necessarily identical to tasks with which the individual is familiar, and a truly *novel problem*.

When practicing chemists first encounter our distinction between an exercise and a problem, they are often tempted to conclude that difference is one of difficulty, or complexity. Our work suggests that problems are neither inherently more difficult or more complex. The only difference between an exercise and a problem is the element of familiarity. Consider the mechanism problem posed in the introduction. It is a routine exercise for a practicing organic chemist, but a novel problem for students who encounter chemistry for the first time.

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

- They are logical sequences of steps.
- They string together in a linear fashion like links on a chain.
- They proceed from the initial information to the solution. (Except, of course, in organic synthesis, where we start from the solution and work back to the initial information.)

These textbook solutions, which are often mirrored by instructors in the classroom, are examples of a phenomenon that has been called "forward-chaining" or "forward-working." As such, they are examples of how routine exercises are worked by individuals with years of experience with similar tasks [4]. Our work has shown, however, that they have little, if any, similarity to the approach successful problem solvers use when they encounter novel problems. Problem solving is a much lengthier and inherently messier process.

The distinction between an exercise and a problem 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 in the course of the semester are routine exercises, and expect that individual to "teach" students for whom the same task is a novel problem.

MODELS OF PROBLEM SOLVING

The long-term goal of our work is the development of a model of problem solving that has two characteristics. First, it must fit our experimental data from interviews with successful problem solvers working on what is, for them, novel problems. 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, the first of which is Polya's stage model [5].

- *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 what we observe in our interviews when we watch successful problem solvers solve novel problems. Consider the following problem, for example, which was used in a study with students who were enrolled in a graduate-level course in mechanistic organic chemistry [6].

Insert Figure 4 here

Almost all the participants immediately proposed a correct mechanism for the first part of the problem, a routine exercise they had seen since their introductory organic courses. When working on the second part, however, the participants who correctly solved the problem talked about their need to "play around with resonance structures" before they could "see" their eventual solutions to what was for them a novel problem. They could not and did not follow the linear process that characterizes Polya's model.

Several other models of problem solving that are logical extensions of Polya's model have been discussed elsewhere [4]. They all have the disadvantage of being inconsistent with the behavior we've observed for successful problem solvers working on novel chemistry problems. Let's therefore turn to a fundamentally different model proposed by Alex Johnstone and co-workers [7]. 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-ElBanna model assumes that when the working-memory capacity of the individual is 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 the presence of specific content knowledge, whether this content knowledge is readily accessible, and on the student's motivation to solve the problem.

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 there is a sharp drop in performance of high-school students when the demand of the problem exceeds capacity. Some students ($\approx$10%) seem to be able to solve problems for which the demand exceeds capacity ($Z > X$), however, because of chunking devices that reduce the demand on working memory.

Let's assume that the Johnstone-ElBanna model is correct when it is applied to situations that meet the six criteria proposed by Tsarpalis, et al. [8]. 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" [9]. Now what? From the perspective of this model, there isn't much we can do to improve student performance in our classes. 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 in mathematics education, Grayson Wheatley proposed an anarachistic model of problem solving that describes what successful problem solvers do when they work on novel problems [3]. In our early work on problem solving in general chemistry courses, we found that this model was consistent with what successful problem solvers did when they encountered novel problems that often had a mathematical component, such as gas law or stoichiometry problems. Recently we have found that this model was the best fit among possible models of problem solving for the results of problem-solving interviews with sophomore organic chemistry students working on tasks that asked them to convert a given starting material into a given target molecule, or to design a set of reactions that would lead to a particular target molecule when the starting material was not specified [10]. This result was interesting because these tasks are fundamentally non-mathematical in nature. Wheatley's anarchistic model consists of the following steps:

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

"Draw a Picture"

There are several stages in the 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 molarity calculation, that drawings are seldom done when we encounter a routine exercise.

In our study of problem solving by graduate students in an organic mechanism course [6] we earlier noted that the individuals who successfully solved the second part of the problem in Figure 4 stated that they needed to "play around with resonance structures" before they could "see" their eventual solutions. In other words, they had to "draw something."

Over the years, several of the first 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.

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"

The steps "try something" and "try something else" may be described as "playing with the problem." Our interviews with beginning students — particularly those who are struggling with a course — have suggested that far too many believe that "trial and error" is not a legitimate strategy for problem solving. This is unfortunate because our work with successful problem solvers suggests that it is among the most powerful strategy these individuals have.

Many examples from our interviews could be used to illustrate the role of "trying something" in the problem solving process, and the reluctance of beginning students to do this. Let's consider an example, however, from our recently completed study of how early-career organic chemists solve organic synthesis problems [11]. A common heuristic used by organic chemists in the design of an organic synthesis is retrosynthetic analysis [12], in which structural features in the products of the reaction are manipulated in a reverse-synthetic sense to deduce the optimum starting materials for the synthesis. One of the participants in our study, who was given the pseudonym William, had been reluctant to apply retrosynthetic analysis when he began the semester-long process of solving the synthesis problem he had been given. Later in the semester, he commented on this reluctance as follows:

> Well, what makes it a little bit easier for me now is to, to not be afraid of just jotting down a bunch of different ways on paper and just, OK, what would happen if I break this bond? OK, well if I break this one then I have this.

What was particularly telling in this interview was William's expressed fear of conducting what is essentially a private exercise that would not be shared with anyone else.

William is not alone in his reluctance to "try something" at the beginning of the problem-solving process. We have repeatedly encountered students in courses at all levels, from general chemistry through graduate school, who were reluctant to "try something," to write something down as a tentative step toward an answer. It is tempting to argue that one source of the resistance to "trying something" is the fact that students seldom see their instructors do this in class because their instructors are working on tasks that are routine exercises for them, not novel problems.

There is also abundant evidence in our data that when successful problem solvers "try something," they tend to take small steps. William, who earlier had expressed his fear of using retrosynthetic analysis further explained how he tried to apply it to his molecule at the beginning of his synthesis project:

> Well, you know, before, the first time I was trying to do it, I wasn't really, I was just trying to do it all in my head the first time and just, just write a retrosynthesis down.

We have found that successful problem solvers tend to take small steps and check to see where they are going rather than taking gigantic leaps, "doing it all in their heads." In a recent study of graduate students and faculty faced with the task of answering questions that ask the individual to deduce the structure of an organic molecule based on a combination of IR and ^{1}H NMR spectra and the molecular formula of the compound, for example, we noted that one of the characteristic differences between those who were judged as "more successful" versus "less successful" was the tendency to write down preliminary results as they deduced fragments of the total structure of the molecule [13].

"Does the Answer Make Sense?"

The penultimate step in Wheatley's model is particularly important. Instructors, struck by the absurdity of some of their students' answers, often conclude that if the students had merely evaluated their answer, they would have caught their own mistake. We have found that beginners seldom test their answer to see if it makes sense, probably for a combination of reasons. First, they have seldom seen their instructors do this when they've watched them work out the solutions in class. Consider the following question, for example.

> What mass of magnesium oxide would be produced when 10.0 grams of magnesium react with excess oxygen?
>
> $$2\ Mg(s) + O_2(g) \rightarrow 2\ MgO(s)$$

How many instructors are likely to calculate the number of grams of oxygen that would be consumed in this reaction and then add this to the mass of the magnesium with which they started in order to "check" the validity of their answer when they get an answer of 16.6 grams of MgO?

Second, they may not have been given the information they would need to check whether their answer "makes sense." Consider an example we have used in seminars and workshops on problem solving, which is based on the first author's experience in a PChem lab almost 40 years ago. He reported the "heat of reaction" in an acid-base neutralization experiment as 13,000 kcal/mol. When the TA's handed back the lab they noted that this value was absurd. (They're right, it should have 13 kcal/mol.) They then said: "You should have known better," to which he responded: "How should I have known better?" At no point in his undergraduate career had anyone given him any basis for predicting what would be a reasonable value for what we now refer to as an enthalpy of reaction.

We have already noted one of the characteristic differences between the behavior of the "more successful" and "less successful" participants in the study of combined spectral interpretation by Cartrette [13]. Another difference that is relevant for this section was the tendency of the "more successful" individuals to check their answer, once they had deduced a molecular structure, against the information in the spectra they were provided.

The Difference Between Exercises and Problems

Our work suggests that Polya's model is better suited for describing what happens when individuals work a routine exercise. They read the question, understand the task, devise a plan, and so on. One of the characteristic tests of whether a task is an exercise for a given individual is to ask: How would one describe the process by which the individual discovered the answer? Exercises are worked in a linear, forward-chaining, rational manner. The anarchistic model of problem solving that we have adopted from the work of Grayson Wheatley suggests that problem solving is cyclic, reflective, and can appear irrational to someone watching the problem solver at work. Experts who watch students struggle with a problem are tempted to intervene; to show the "correct" way of obtaining the answer. This may make the expert feel good, but it doesn't necessarily help the individual struggling with the problem for the first time.

REPRESENTATIONS AND REPRESENTATIONAL SYSTEMS

Our first hint into the role that representations and representational systems play in problem solving in chemistry came from a study in which we looked at students' answers to the following question [14].

> Predict the product of the following reaction:
>
> $$PhCOOH + SOCl_2 \rightarrow$$

A couple of typical incorrect answers are given below.

$$PhCOOH + SOCl_2 \rightarrow PhCl + SO_2 + HCl$$

$$PhCOOH + SOCl_2 \rightarrow PhCOOCl + SO_2 + HCl$$

When people who teach general chemistry look at these equations, they often note that they are not balanced. That doesn't bother us because organic chemists seldom worry about the mundane details of writing balanced equations. What bothers us is the fact that these equations are absurd; there is no way to go from the starting materials to the products of these equations by making and breaking of bonds.

When we did this experiment, we noted that some of the students who answered this question correctly wrote a symbolic representation for $PhCO_2H$ in the margin of the exam paper. It seldom looked as regular as the symbolic representation in Figure 5.

Insert Figure 5 here

Sometimes the ring looked as if it had six "bumps" corresponding to the six carbons of a benzene ring, often it did not. The ring sometimes contained one double bond, sometimes two, sometimes three. Sometimes it didn't contain any double bonds. But the -CO_2H portion was always clearly written.

Some would argue that the "$PhCO_2H$" with which the starting material was presented on the exam is a "symbolic" representation. We'd like to argue that it *can be* a symbolic representation, but it often is not. For many students, particularly those who struggle with organic chemistry, it is a verbal/linguistic representation that consists of letters and numbers that aren't symbols because they don't symbolize anything.

Interviews with students struggling with organic chemistry have lead us to conclude that there is a fundamental difference between what the instructor writes on the blackboard and what students write in their notebooks, in spite of the fact that one seems to be a direct copy of the other [15]. The instructor writes *symbols*, which represent a physical reality. All too often, students write *letters and numbers and lines*, which aren't symbols because they have no physical meaning to them. Interviews suggest that it is the students who are trapped in verbal/linguistic representation systems who are most likely to write the equation in Figure 6 to represent attack by a Grignard reagent on a ketone.

Insert Figure 6 here

It isn't until the letters, lines and numbers in this equation become symbols that this answer becomes wrong. One can draw as many lines to a "C" as one wants; it is only bonds to a carbon atom that are limited to four.

ACCULTURATION AND PROBLEM SOLVING

If one accepts the notion that are fundamental differences in the way in which individuals approach routine exercises and novel problems, one might conclude that working examples in class that reflect how we, as instructors, solve routine exercises is not likely to help students become better problem solvers. We have reason to believe, however, based on our experiences teaching general chemistry and, more recently, physical chemistry that students' problem-solving ability can be improved by doing two things. First, exposing them to the anarchistic model of problem solving as an example of how successful problem solvers approach novel problems. And second, using this model explicitly in class over and over again during the course of the semester to demonstrate how novel problems can be worked.

Recent results from our work has suggested another way to help students improve their problem-solving ability. Greeno and Hall [16] have suggested that "learning to construct and interpret representations involves learning to participate in the complex practices of communication and reasoning in which the representations are used." In our recently completed study on how graduate students learn to design organic syntheses, we used an ethnomethodological approach to understand the experiences of newcomers to the culture of organic chemistry [11]. The participants in that study described how their exposure to the culture of organic chemistry that occurs in the research lab, during group meetings, and during conversations about research among their peers helped them learn several useful problem-solving strategies. The data from that study suggest that the process of "acculturation" that occurs as graduate students change some of their behaviors from those of "students" to those of "practicing chemists" helped the participants in this study become better at solving problems within the domain of synthetic organic chemistry.

In traditional lectures, instructors tend to focus on concepts and principles of their field and expect that students will be able to apply this material to problem-solving activities on their own. This might result from the fact that science instructors frequently approach the task of helping students learn science as if the students were learning a new language. However, concentrating on the language of science — the concepts and principles of organic chemistry or physical chemistry, for example — does not necessarily convey the tools and methods that students need to create their own "meaningful phrases" in the language of chemistry, i.e., successfully solve problems. We believe that we can help our students develop into more successful problem solvers by teaching science as a culture, since culture encompasses so much more than just language. The notion of culture combines language with embedded meaning and with the tools people in their respective communities use.

This stance has consequences for both science education research and for teaching. First, we believe that additional ethnomethodological studies in specific domains of science need to be conducted. The problem-solving methods that practitioners of a discipline employ are so ingrained in their daily activities that they are implicitly embedded in the culture of the practice and often cannot be explicitly stated by any one of the individuals who use them. Our proposed line of research is likely to bring to surface and therefore make explicit the functional problem solving heuristics that practitioners use in their respective disciplines.

In addition to suggesting new avenues for research, treating science as a culture has implications for teaching as well. Traditional classroom instruction has been based on the implicit assumption that the epistemological development of learners does not depend on the domain in which learning takes place. Consider, for example, Perry's [17] now-famous model of the intellectual development of college-aged males, which is supposed to be domain independent. More recent work [18], however, suggests the need for the reinterpretation of epistemological development from a domain-specific model.

In his work on the learning of quantum mechanics, Gardner [19] showed that when students were exposed to the overarching framework of physical chemistry they had a significantly smoother experience learning quantum mechanics. Our work on problem solving in organic chemistry supports the notion of domain-specific epistemological development. For example, the first-year graduate students in our study of organic synthesis were good, successful students who had read vast volumes of textbook materials by the time they had begun graduate school. In the first half of their graduate-level organic synthesis course, however, it was obvious that they didn't know how to read journal articles that described the results of synthetic organic chemistry research. They weren't paying attention to factors such as yield and stereoselectivity because they didn't know that they were supposed to pay attention to such items. Success in the organic synthesis problem-solving venture only came after these participants used their exposure to the culture of practicing organic chemists to construct meaning and, thereby, learn how to use some of the heuristics of organic chemistry.

It is not reasonable or even desirable to immerse all the undergraduates enrolled in introductory organic chemistry in the culture of organic chemistry in the form of extensive exposure to the research laboratory. Since we cannot take the students to the culture; we must bring the culture to them. One way this can be achieved is by giving the students an overarching framework of the field that is introduced at the beginning of the course and reinforced throughout its duration.

Not only would giving students a sense for the framework of organic chemistry increase the likelihood that they will know what to study, it would also increase the likelihood that the student would know what to transfer from one situation to the next. Transfer, of course, plays a particularly important role in organic chemistry. Practicing organic chemists would cite the addition reactions of alkenes, for example, as evidence that alkynes should undergo similar reactions. Having undergone the process of acculturation into the culture of organic chemistry, they would use the overarching framework of the field to justify such reactions if they found that they did indeed occur, attributing the reaction to the nucleophilicity of the π electrons in the carbon-carbon triple bond.

In addition to providing an overarching framework of the discipline, teaching science as a culture requires more emphasis on the explicit instruction in the way tools are used within one the specific domains of science. In the context of organic chemistry, examples of tools would be the arrow-pushing formalism and retrosynthetic analysis [11]. It is difficult to imagine an organic chemistry course, from the sophomore-level to graduate school, that does not use the arrow-pushing formalism. Recent work in our group suggests, however, that instructors significantly overestimate the extent to which students understand what their are doing when they utilize this formalism in class [20]. As one of the primary tools that help transform the language of organic chemistry into the culture of practicing organic chemists, it is essential to consistently remind students of the meaning of each step when presented with mechanisms, i.e., the significance of each arrow and how those arrows correlate with the chemical principles involved in that step. An equivalent amount of instructional time needs to be spent on retrosynthetic analysis and how it can be used to simplify a synthesis problem.

We believe that classes that focus on the concepts and principles of chemistry should be supplemented with instruction where the explicit application of that theory to problem solving is shown. One of the prime ways to implement this is by using Peer-Led Team Learning (PLTL) approaches [21]. Not only can PLTL workshops provide the necessary instruction for its students, but the workshop environment incorporates many of the useful characteristics of the culture of organic chemistry.

Defining heuristics that practitioners use and teaching science as a culture are only the first steps in helping our students develop as problem solvers. Ultimately, the instructor needs to incorporate these ideas into solving *problems* during lectures. This means that reasonable pathways that don't lead to the "correct" answer need to be explored and analyzed. Only by doing this can we help students construct a more realistic notion of how they go about the process of solving novel problems when they arise.

REFERENCES

[1] G. M. Bodner, *University Chemistry Education*, **2003**, *7*, 1-9.
[2] Hayes, J. *The Complete Problem Solver*. The Franklin Institute: Philadelphia, 1980.
[3] Wheatley, G. H. Problem solving in school mathematics. MEPS Technical Report 84.01, School Mathematics and Science Center, Purdue University, West Lafayette, IN, 1984.
[4] Bodner, G. M.; Herron, J. D. "Problem Solving in Chemistry," in *Chemical Education: Research-Based Practice*, J. K. Gilbert, Ed., Kluwer Academic Publishers, 2002.
[5] Polya, G. *How to Solve It: A New Aspect of Mathematical Method*. Princeton University Press: Princeton, NJ, 1945.
[6] Bhattacharyya, G.; Bodner, G. M. *Journal of Chemical Education*, **2005**, *in press.*
[7] Johnstone, A. H.; El-Banna, H. *Education in Chemistry*, **1986**, *23*, 80-84.
[8] Tsaparlis, G., Kousathana, M., Niaz, M. *Science Education*, **1998**, *82*, 437-454.
[9] Niaz, M. *Journal of Chemical Education*, **1987**, *64*(6), 502-505.
[10] Calimsiz, S. How Undergraduates Solve Organic Synthesis Problems: A Problem Solving Model Approach, Unpublished Master's Thesis, Purdue University, May, 2003.
[11] Bhattacharyya, G. A Recovering Organic Chemist's Attempts at Self-Realization: How Students Learn to Solve Organic Synthesis Problems, Unpublished Ph.D. Dissertation, Purdue University, December, 2004.
[12] Corey, E. J.; Cheng, X. *The Logic of Chemical Synthesis*. J. Wiley & Sons: New York, 1989.
[13] Cartrette, D. Using spectral analysis to probe the continuum of problem solving ability among practicing organic chemists. Unpublished Ph.D. Dissertation, Purdue University, December, 2003.
[14] Pribyl, J. R.; Bodner, G. M. *Journal of Research in Science Teaching*, **1987**, *24*, 229-240.
[15] Bodner, G. M.; Domin, D.S. *University Chemistry Education*, **2000**, *4*, 24-30.
[16] Greeno, J.; Hall, R. *Phi Delta Kappan,* **1997**, *78*, 361-367.
[17] Perry, W. *Forms of intellectual and ethical development in the college years: A scheme.* Holt, Rinehart, & Winston: New York, 1970.
[18] Carey, S.; Smith, C. *Educational Psychologist*, **1993**, 28, 235-251.
[19] Gardner, D. Learning in quantum mechanics. Unpublished Ph.D. Dissertation, Purdue University, December, 2002.

[20] Ferguson, R. Understanding Arrow-Pushing Formalism from a Student's Perspective. Unpublished Ph.D. Dissertation, Purdue University, August, 2003.
[21] Tien, L.; Roth, V.; Kampmeier, J. *Journal of Research in Science Teaching*, **2002**, *39*, 606-632.

Figure 1

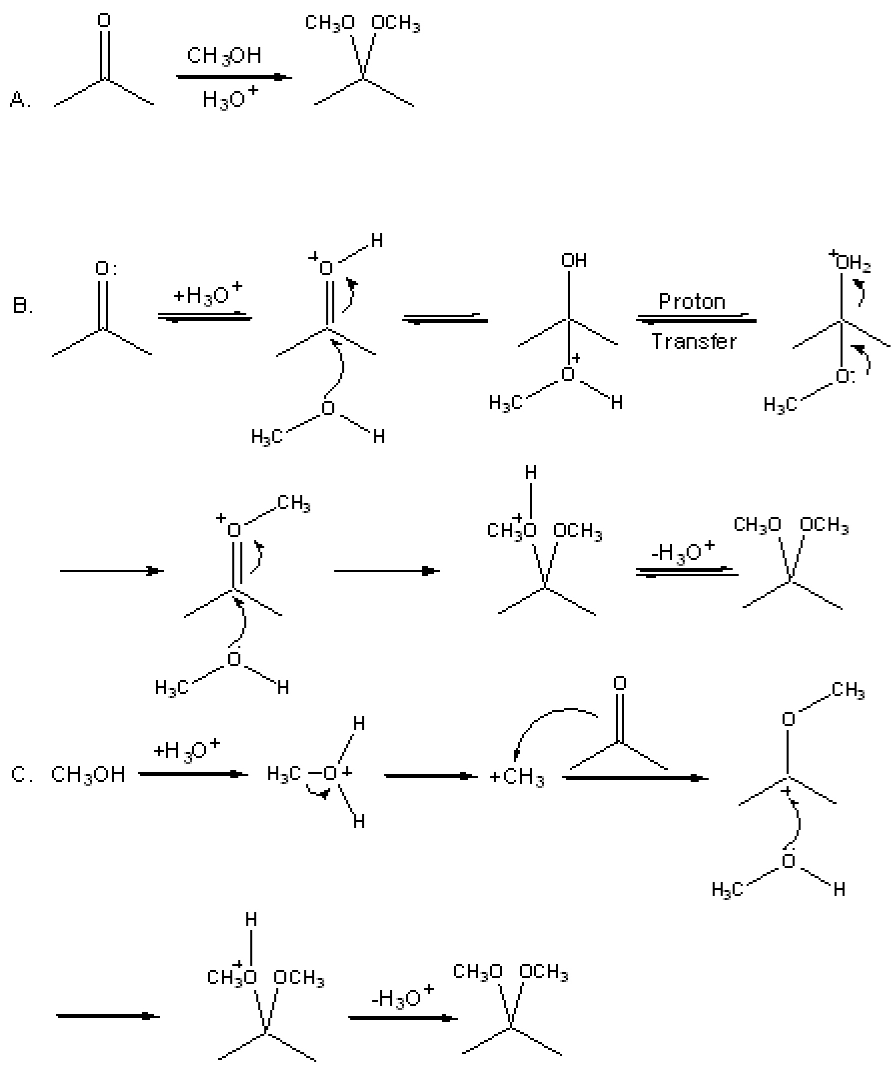

Figure 1: A sample mechanism problem that might be given on a sophomore-level organic chemistry exam. (A) The acetal formation reaction for which the mechanism should be written. (B) An answer that would be accepted by chemists as the “correct” mechanism for the reaction. (C) A common, but incorrect, answer that students give.

Figure 2

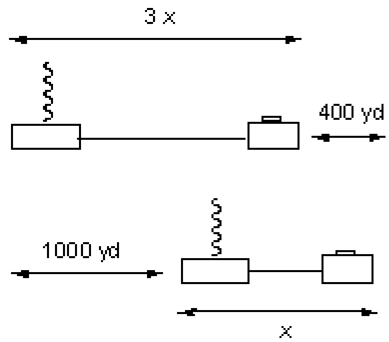

Figure 2: A representation of the drawings constructed by virtually every practicing chemist to whom the two-trains problem is given.

Figure 3

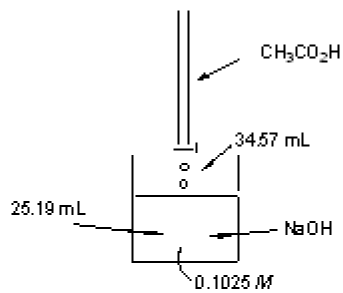

Figure 3: An example of a drawing that *could* be constructed from the information in a standard titration problem.

Figure 4

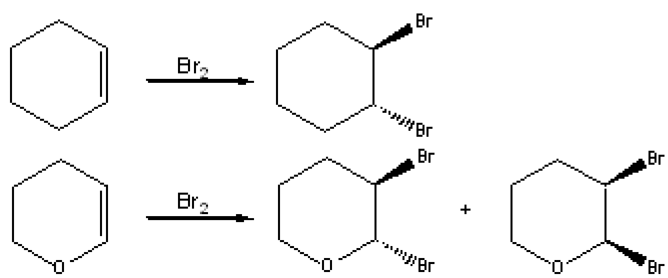

Figure 4: A two-part mechanism problem that was used in a study of the use of organic reaction mechanisms by graduate students. The participants were asked to use mechanisms to explain the difference between the products of the two reactions.

Figure 5

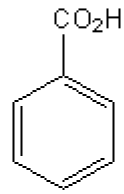

Figure 5: Students who successfully predicted the product of the reaction between "$PhCO_2H$" and $SOCl_2$ often drew a symbolic representation that contained at least some of the information in this figure, although it seldom looked this regular.

Figure 6:

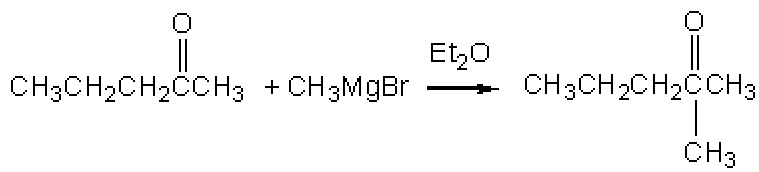

Figure 6: An example of the kind of answer often given by students who are locked in a verbal/linguistic representation systems when asked to predict the product of a Grignard reaction.