

Conclusions.

This research project produced some findings that can potentially affect many aspects of the educational enterprise related to teaching and learning problem-solving skills. The summary of these findings is presented below.

Educational research.

This study, to my knowledge, is the first of its kind in several respects. The innovative character of my methodology has been discussed in great detail in Chapter 2; however, as I am pondering the outcome of my study, a reasonable question to ask is: “Just how useful is this new methodology in constructing new knowledge?” The outcome of the project has shown that it is, in fact, very promising.

A new theoretical framework (BARK) was developed as a foundation for my methodology. This new framework is well-grounded in the existing theories but takes them further by providing a useful way to describe the cognitive skills that determine one’s ability to solve *challenging* problems.

The tasks developed for the study also turned out to be appropriate for my purposes. The multiple-choice test (Mechanics Diagnostic Test) proved to be a reasonably good diagnostic tool for predicting one’s ability to solve problems. I view this fact as a great success, since most well-known diagnostic tests used in physics education (such as Force Concept Inventory, Mechanics Baseline Test, etc.) took years of development and fine-tuning by the research teams with resources far greater than mine. The open-response tasks appeared to be have the right level of difficulty, providing serious challenge for most students but being manageable – my previous experience as a problem-writer turned out to be very helpful.

In using the newly developed interactive web platform CyberTutor, I managed to conduct a problem-solving study with more than 100 high-school participants from several geographic areas of the country. Nothing like this could have been manageable just several years ago – and the advantage of using larger numbers of participants is fairly obvious. In the past, researchers attempted to generalize from studying the problem-solving processes of as few as two or three participants, and such results are bound to have serious limitations. Using large numbers of participants has at least two clear advantages. First, the researcher is able to apply statistical techniques to the data, thus expanding the ways in which the evidence is analyzed. Second, the fact that the students come from many different classes – not from one class, as has often been the case in the previous studies – makes the results more valid, eliminating the possible influence due to a particular teaching style or a particular set of tasks to which the participants may have been exposed.

The use of CyberTutor not only enabled me to administer the tasks to a very large sample of students in an efficient and consistent manner; it also *provided new ways to collect information about the students’ ways of thinking*. By emulating Socratic dialog through pre-programmed hints, CyberTutor allowed me to obtain far more detailed and better structured data than a simple right/wrong feedback that most “interactive” computer programs provide.

From my review of prior research and from the results of this study, several methodology-related conclusions emerge:

- Any researcher investigating problem solving must pay serious attention to the tasks offered to the participants. In the past, “problem-solving skills” were usually tested by offering the participants relatively trivial tasks¹, the tasks that I would classify as exercises, not problems.
- Larger and more diverse samples of participants provide data of greater validity and can be generalized with greater confidence. Modern technologies, including the Internet, provide new great opportunities in conducting research with such samples; these opportunities should be fully utilized.
- If the researcher’s goal is to observe the differences between “good” and “bad” problem-solvers, it is more useful to compare the participants with *similar* backgrounds, not those with vastly different amounts of experience in the field.
If the researcher chooses to compare the problem-solving behavior of such “differently-experienced” individuals, different sets of tasks may have to be offered to them.
- It is useful to combine qualitative and quantitative methods of research in studying problem-solving. The insights provided by these methods appear to complement each other rather well,

¹ For examples of such studies, see Differences between experts and novices in problem solving in Chapter 1.

allowing the researcher to cross-check some conclusions and to obtain a richer and more valid array of data.

- It is very helpful to use a sample of students with very similar backgrounds in physics; that way, the researcher can be sure that the tasks will be, objectively, of about the same degree of familiarity for all participants. If some participants are far more familiar with particular tasks than are others, the cognitive skills involved in solving will be very different, and any comparison drawn from such studies, will have limited value for advancing problem-solving research.
- Using pre-programmed hints tremendously improves the quality of information that can be obtained from the participants' interactions with computers. Researchers often ask the participants to explain what made the problems difficult and what steps they took to overcome these difficulties. Such a method, while useful (I collected that information too, from the written comments), is hardly a perfect probing tool. Self-analysis is difficult, especially for an inexperienced individual: it is not easy for the students, I found, to retrace their thinking – especially given the non-sequential way in which problem-solving processes tend to develop once the tasks get challenging. The use of hints, and their *recordable* effect on the solution process, provides additional perspective and, if developed further, may allow researchers to conduct high-quality *virtual interviews* with very large numbers of participants. Not only will the information obtained in such a way be vast – it will also be easily manageable, owing to its computer recorded format.

This study can be helpful to future researchers not only from the methodological perspective. Several theoretical aspects of existing problem-solving research have been challenged, reformulated and developed beyond the state found in existing literature:

- *Rigid knowledge* and *bisociation*, as shown by my results, do appear to be distinct components of the problem-solving ability, and this perspective will hopefully advance the ways in which problem solving is taught and learned.
- Clarifying the distinctions between the “*expert – novice*” and the “*good solver – poor solver*” dichotomies, I offer a way to narrow the focus of one's study and to explore only the particular kinds of expertise in which the researcher is interested.
- Using my operational definitions of *problems* and *exercises*, I provide a unifying way in which to view the existing literature on problem-solving research and to formulate new research questions. In particular, I hope to help future researchers focus on developing the problem-solving skills separately from the exercise-solving ones.

Teacher training.

Given the large amount of evidence collected by both researchers and practitioners, it is not entirely clear to what extent teaching general problem-solving skills *is* possible. How so?

Usually, the teachers, from secondary-school to college level, teach their students to solve *exercises*, day in and day out; meanwhile, *some* of their students also show *talent* for solving challenging tasks – the tasks that I define as *problems*. I have been fortunate to work with students of a very wide range of abilities, from “regular C” students to those who won gold medals on the International Olympiads and could solve problems far better than their teachers. Good teaching was, I am sure, important in inspiring these students' interest and in providing them with necessary *rigid knowledge*. Meanwhile, it is not clear that their *bisociation* skills which, in my view, are crucial to one's ability to solve problems, were the direct result of teaching.

All attempts to teach problem-solving consistently to large groups of students have had very limited success – the evidence supporting this statement is discussed in detail in Chapters 1 and 2. Still, most teachers of physics would view teaching problem-solving skills as a high priority. As a teacher and as a researcher, I believe that problem-solving skills can be taught – with at least a greater degree of success than has been achieved so far. However, given the general failure of educational practitioners in that regard (*exercise-solving* skills are different; they *have* been taught well to many students), it seems obvious that *practitioners need training* in effective teaching of such skills.

My research provides guidance for designing such training programs. The main idea of such training, in my opinion, is this: *the pedagogy of teacher training must mirror the pedagogy that the teachers are expected*

to practice. The name for such pedagogy has been proposed in the title of this chapter: **BITE² – Bisociation-Intensive Teacher Education**.

Here are the basic principles of this pedagogy:

- Teachers need to *experience* problem-solving during their training; that is, they need to be offered challenging tasks to solve; the tasks that make the teachers struggle just like the students would struggle with *their* tasks.
- In the process of solving, the teachers should be *explicitly taught the bisociation skills* and solving strategies that are found to be useful; the list of such strategies will follow in the Classroom Teaching section of this chapter.
- Teachers, just like students, *must find a good degree of success* in solving the challenging tasks that they will be offered; the teacher who has rediscovered (or, maybe, experienced for the first time!) a true joy arising from cracking a tough problem is a teacher who is more likely to convey *the joy and the hope of its anticipation* to his or her students.
- Since trying alternative approaches is one of the key strategies in solving difficult problems, the teachers should be trained to consistently *show the students that any problem can be solved in multiple ways*, to stress the idea that such multiple ways do exist and do deserve to be considered.
- Last but not least: the real problem-solving processes are *non-linear*. The solver often has to come back to the previous stages of the solution to look for new relationships, to question previous assumptions, etc. However, a conventional way of demonstrating “problem-solving techniques” in class by teachers is showing “linear solutions”, when the teacher knows exactly what to do and shows no signs of “mental disequilibrium” (Lawson & Wollman, 1975). In such examples, each step is neatly followed by the next logical step, and there are no mistakes made – ever. There is no self-monitoring involved, for the teacher knows what the next step is. Usually, the students watch in awe and say thing like “*Now that you showed us, it is all so clear. But how was I to know to take that step here?*” Such teaching practice gives the students illusions that, if they study well, all their solution processes should be flawless and straightforward – which can only be true for the most straightforward exercises. To correct this epistemological misconception, the teachers must be trained to demonstrate the process of solving difficult and *not* pre-solved problems in front of the class, so that the students can see for themselves the processes of questioning oneself, dealing with confusion, discarding the wrong assumptions, trying alternative approaches, self-monitoring techniques, etc.

In my view, *problem-solving training should be an integral part of physics teachers’ training in general*. Not only because teaching problem-solving skills is commonly viewed as an important part of formal education – but also because *problem-solving is what scientists do*. The nature of a scientist’s work is noticing the connections that have not been noticed before and discovering the hidden features of the unfamiliar objects and phenomena. Thus, a student who appreciates the value and the joy of problem-solving is more likely to understand the nature of the intellectual demands and the career options in the area of science.

Curriculum development.

One of my key propositions, well supported by the results of the student interviews, is that *true problem-solving skills are not engaged unless they are called for*. Many studies in which the students soundly rejected seemingly useful problem-solving skills are a sound proof of that³. Research suggests that individuals pursue the least cognitively demanding ways to solve the tasks offered to them⁴ – thus a teacher should not expect a student who possesses decent problem-solving skills to deploy them on a typical textbook end-of-chapter “problem.”

Therefore, successful teaching of problem-solving skills necessitates *both* well-trained teachers who can effectively present useful problem-solving strategies to their students and a good selection of challenging, stimulating tasks that would *force* the students to use these effective strategies. The analysis of the current textbooks, presented in Chapter 1, clearly shows that the packages of problems currently available to most physics students, both at high-school and college level, are far away from being satisfactory. Currently, the

² Once again, marketing considerations took the better of me...

³ See Chapter 1, Teaching Problem-Solving section for more detailed treatment.

⁴ This idea is well-developed in the literature; see discussion in Chapter 2 for details and references.

quality and the style of the problem content is about the same for most textbooks, which tend to compete on the quality of the “explanations.” Research dictates, however, that effective teaching of problem-solving skills requires good tasks much more than good textbook explanations. The priority of the textbook authors and curriculum developers in general needs to be shifted from *chapters* to *end-of-chapter tasks* – as it stands now, their quality is very low, and the expression “end-of-chapter task”, in a conversation of physics educators, sounds almost derogatory.

The physics education community has long developed a coherent view of what constitutes a good problem⁵. I suggest that publishers and educational foundations heed this view and dedicate substantial resources to developing textbooks filled with high-quality tasks, not just routine exercises but also challenging problems of different style, degree of difficulty, etc.

Classroom teaching.

Assuming that well-trained teachers and good problem packages are available to the students of the future, there is a good chance that they will be taught problem-solving skills more successfully. The skills in question are mostly bisociation-related and, judging from my study, they may, in fact, be teachable. I would like to point out again that, according to my results, better problem-solvers do not use solving strategies that are drastically different from the ones employed by weaker problem-solvers. Rather, more successful problem solving is associated with more efficient, consistent and creative use of the same array of solving approaches. Thus, the explicit teaching of such approaches may hold the key to substantial improvements in making the students more successful problem solvers.

I would like to offer a list of “points” that need to be made during the teaching of problem solving. This list is by no means complete – nor can it possibly be. It is not static, either: I fully expect it to be someday modified by future research in this area. Finally, given the quantitative nature of the tasks used in my study, it is geared toward quantitative problems, which limits it somewhat.

Still, I feel fairly confident that the following ideas should be conveyed to – and experienced by – the students who are to become stronger problem-solvers:

- Before you approach a problem, you must *assume that it will be challenging*. You may have to spend quite a bit of time solving it; your initial approach may well fail. This is *normal*, and you should not feel discouraged.
- Before writing *anything*, ponder the problem, visualize the situation, think about the knowledge that you may find useful in solving it.
- If the situation looks totally unfamiliar, try to look for *at least some* familiar features within it; your experience *should* be sufficient for finding them.
- Ideally, you should have a path to the answer broadly mapped out before you proceed to look for specific equations. Even if this path eventually leads nowhere, it helps to have it. If such “advance mapping” seems too hard to do, at least use short-term planning. Whenever you consider another step, ask yourself: What is next? Why am I trying to find this? Will this help?
- If the problem is quantitative in nature, make sure that you understand the way in which the desired answer can be expressed mathematically⁶.
- Once the unknown is expressed mathematically, create an equation or, depending on the need, a system of equations that can be solved to find the unknown. *This is the key part of the process: expect to have difficulties here*. You will have to use the most general physics principles as well as the very specific inner features of the problem situation to elicit the necessary equations.
- Any interesting, non-trivial relationship that you find may be turned into an equation and used in your solution. On the other hand, even a generally correct equation may not be applicable to the specific problem situation or it may not be relevant to the solution. *Question the relevance of any equation* that you enter into your system of equations.
- If you cannot assemble the necessary equations, *consider an alternative approach*: search for the inner features of the situation that may be more helpful than the ones you considered before; try to

⁵ See Chapter 1, What Makes a Good Problem?

⁶ In my tasks, formalizing the unknown was not an issue. However, many challenging problems pose the question in ways that are hiding the mathematical nature of the unknown, e.g. “... Given this information, will Jimmy catch the train?”

apply the concepts that are applicable but that you have not yet considered. Of course, it helps to know your basics.

- If you feel that you have enough equations to start solving, proceed with caution. Pause to inspect your intermediate results – do they make sense? If not, is it a technical error or a wrong assumption made when the equation was first constructed?
- Do not be afraid to introduce new unknowns. Either you will be able to solve for them or they will not be a part of the final answer⁷.
- If you discover an error or hit a dead end, please remember: this is normal. Retrace your solution to the initial equations. Are they all correct? Are they all useful? Can some of them be replaced? To replace them, you will need to discover something else about the problem situation.
- Be patient. Try again. If that fails – try again. If the problem situation is too complex to grasp, think about ways to simplify it. Solve a simpler sub-problem, then come back to the original one.
- If you finally assemble a solvable system of equations and solve it, examine the answer. If the answer is numeric, is the value obtained realistic? If algebraic, is it dimensionally correct? Does it hold true for extreme cases? Are the initial assumptions reasonable?
- Reflect. What does the answer mean? What is it useful for? Could this problem be solved in a different way?
- You just solved a tough problem. You went through disappointments, dead ends, frustration – but you solved it. Wasn't it *fun*?

Future research.

Now that the project is complete, I can say that the theoretical framework developed from analysis of prior research was very helpful in formulating the methodology of my study; the project, especially given its inevitable constraints, was a definite success. However, as I hope to pursue this line of research further, I would like to propose several ways in which my research methodology can be enhanced.

Development of research methodology.

- I would like to recruit even larger numbers of students, so that each of the various *sub-samples* – such as females, juniors, students of particular ethnicity, students with particular background in physics and/or mathematics – has a substantial number of students allowing me to apply statistical techniques to the sub-samples in a meaningful way.
- I would like to develop Mechanics Diagnostic Test further until it contains about 30-35 questions, equally and *sharply* split between the R-type and the B-type. This will bring the number of MDT questions in line with most well-known diagnostic tools currently in use.
- I also expect to improve the quality of the open-response tasks. The limitations imposed by this study did not allow for extensive pilot-testing of the problems and the hints. In the future, it should be possible to create and pre-test a greater number of the new tasks, increasing both the quantity and the quality of information that can be collected from exploring the student interactions with such tasks.
- Ideally, it would be great to have all participants solve problems in a controlled, supervised environment. It is possible if the advance agreement with the instructors is reached, and there is time allocated for the students to solve the problems in the classroom. This way, student effort will be more consistent, and the results may be more consistent, too. I think that such an arrangement is realistic if the instructors are properly recruited and, perhaps, compensated for their time and effort.
- Student comments can be solicited in a more structured way: for instance, a “comment window” can open as the student submits the final answer to a problem; that window can contain specific questions, making the comments, in effect, “virtual interviews” with results that can be easily tabulated and analyzed.

⁷ If neither seems to be the case, then the problem may not be stated correctly, of course.

Possible focus of future studies.

I expect to use this improved methodology to advance this research in several ways:

- to conduct larger-scale studies with high-school and college students, in an attempt to obtain more solid experimental evidence that would help to confirm and, perhaps, modify my current theoretical framework of problem-solving processes;
- to explore the gender-associated differences in problem-solving approaches and the factors associated with success in problem solving. While my study is not conclusive in that regard, there are some hypotheses to be tested; I find this line of research to be quite promising;
- to conduct similar studies in subjects other than physics, especially mathematics. Mathematics seems an important research field, for several reasons. First, it is a discipline centered around problem solving. Second, compared to physics, mathematics is studied by a much larger population of students, making recruitment easier. Finally, mathematics is studied at different ages, and such research may help analyze the developmental influences on individuals' problem-solving abilities. Together, research results in physics, mathematics and, maybe, chemistry may help analyze the most general problem-solving abilities, independent of the academic field.
- I also hope to conduct studies focusing on problems rather than on students. The list of possible research questions includes:
 - What makes a given problem hard?
 - What knowledge or strategy helps solve it?
 - Which types of problems are “bisociation-heavy” and which ones tend to draw more on the students' rigid knowledge?
 - How well are the findings of this study applicable to *qualitative* problems?

By having large numbers of students solve a particular problem in a supervised environment and analyzing the solving processes, new insights about creating pedagogically useful tasks *and* new teaching methods may be gained. As mentioned before, successful teaching of problem-solving skills requires both.

The final word.

This study was designed over the course of several years, and many more years of my experience as a teacher, problem writer and, since recently, a researcher played a role in preparing me for carrying it out. This project is a culmination of my continuing effort to make a positive difference in science education. The study produced some new insights that are fairly reliable, pleasantly non-trivial and, I believe, applicable in many different ways. If my work contributes to advancing the current state of teaching, especially the problem-solving aspect of it, I will be most satisfied.