PUZZLE-BASED LEARNING:
An introduction to critical thinking, mathematics, and problem solving

Zbigniew Michalewicz
Matthew Michalewicz

Published by Hybrid Publishers

Melbourne Victoria Australia

First published 2008

National Library of Australia Cataloguing-in-Publication data:

Michalewicz, Zbigniew.

Puzzle-based learning : introduction to critical thinking, mathematics, and problem solving.

ISBN 9781876462635.

Puzzles. Educational games. Problem solving. Skills. Critical thinking

Michalewicz, Matthew

793.73

Foreword

Google is a company that is renowned for its love of puzzles. We solve puzzles to relax, we subject interview candidates to them, and we even run puzzle competitions. “Googlers” are not alone as people around the world have been fascinated by puzzles for thousands for years.

Solving puzzles is more that mental aerobics though. Like philosophers and mathematicians before them, Zbigniew and Matthew Michalewicz have recognized the pedagogical power that lies in solving puzzles. This book is chock-a-block with interesting puzzles and their solutions, lavishly and wittingly explained. Any reader with a basic knowledge of mathematics plus an ounce of curiosity will find this book enjoyable to read. But the Michalewiczs go further in presenting the problem-solving strategies and principles underlying puzzle solving, and in doing demonstrate the power of puzzle-based learning; that learning problem solving can be fun!

In doing so they have given us a tremendous book about problem solving that is both educational and entertaining at the same time, and one that I hope will be incorporated into problem-solving curricula around the world.

Alan Noble
Engineering Director, Google
Sydney, Australia

To György Pólya and Martin Gardner, who paved the way, and to our Families,
for their patience and understanding during this project.

Z. M. & M. M.

Preface

"Elementary," said he.

The Crooked Man

What is missing in most curricula – from elementary school all the way through to university education – is coursework focused on the development of problem-solving skills. Most students never learn how to think about solving problems. Throughout their education, they are constrained to concentrate on specific questions at the back of textbooks. So, without much thinking, they apply the material from each chapter to solve a few problems given at the end of each chapter (why else would a problem be at the end of the chapter?). With this type of approach to "problem solving," it is not surprising that students are ill prepared for framing and addressing real-world problems. When they finally enter the real world, they suddenly find that problems do not come with instructions or textbooks.

Although many educators are interested in teaching "thinking skills" rather than "teaching information and content," the fact remains that young people often have serious difficulties in independent thinking (or problem-solving skills) regardless of the nature of a problem. As Alex Fisher wrote in his book, *Critical Thinking*: "… though many teachers would claim to teach their students 'how to think', most would say that they do this indirectly or implicitly in the course of teaching the content which belongs to their special subject. Increasingly, educators have come to doubt the effectiveness of teaching 'thinking skills' in this way, because most students simply do not pick up the thinking skills in question." This approach has dominated the educational arena – whether in history, physics, geography, or any other subject – almost ensuring that students never learn how to think about solving problems in general.

Over the past few decades, various people and organizations have attempted to address this educational gap by teaching "thinking skills" based on some structure (e.g., critical thinking, constructive thinking, creative thinking, parallel thinking, vertical thinking, lateral thinking, confrontational and adversarial thinking). However, all these approaches are characterized by a departure from mathematics as they concentrate more on "talking about problems" rather than "solving problems." It is our view that the lack of problem-solving skills in general is the consequence of decreasing levels of mathematical sophistication in modern societies.

Hence, we believe that a different approach is needed. To address this gap in the educational curriculum, we have created a new course (based on this book) that focuses on getting students to *think* about framing and solving *unstructured* problems (those that are not encountered at the end of some textbook chapter …). The idea is to increase the student's mathematical awareness and problem-solving skills by discussing a variety of *puzzles*. In other words, we believe that the course should be based on the best traditions introduced by Gyorgy Polya[1] and Martin Gardner[2] during the last 60 years. In one of our favorite books, *Entertaining Mathematical Puzzles*, Martin Gardner wrote:

> Perhaps in playing with these puzzles you will discover that mathematics is more delightful than you expected. Perhaps this will make you want to study the subject in earnest, or less hesitant about taking up the study of a science for which a knowledge of advanced mathematics will eventually be required.

Many other mathematicians have expressed similar views. For example, Peter Winkler in his book *Mathematical Puzzles: A Connoisseur's Collection* wrote: "I have a feeling that understanding and appreciating puzzles, even those with one-of-a-kind solutions, is good for you."

As a matter of fact, the puzzle-based learning approach has a much longer tradition than just 60 years. The first mathematical puzzles were found in Sumerian texts that date back to around 2,500 BC! Yet the best evidence of the puzzle-based learning approach can be found in the works of Alcuin, an English scholar born around AD 732 whose main work was *Problems to Sharpen the Young* – a text which included over 50 puzzles. Some twelve hundred years later, one of his puzzles is still used in countless artificial intelligence textbooks![3]

So, what is "a puzzle"? Of course, it is difficult to give a universal definition as sometimes the difference is not clear between a puzzle and a real problem. However, in this text we concentrate on educational puzzles that support problem-solving skills and creative thinking. These educational puzzles satisfy most of the following criteria (also see the preface in Peter Winkler's book *Mathematical Puzzles: A Connoisseur's Collection*):

1. *Generality*: Educational puzzles should explain some universal mathematical problem-solving principles. This is of key importance. Most people agree that problem solving

---

[1] Gyorgy Pólya was born in Budapest on 13 December 1887. For most of his career in the United States he was a professor of mathematics at Stanford University. He worked on a great variety of mathematical topics, including series, number theory, combinatorics, and probability. In his later days, Gyorgy Pólya spent considerable effort on trying to characterize the general methods that people use to solve problems, and to describe how problem-solving should be taught and learned.

[2] Martin Gardner was born in Tulsa, Oklahoma on 21 October 1914. He is one of the most beloved personalities in the areas of recreational mathematics, magic, and puzzles. The influence of his work is immeasurable. Martin Gardner is the author of more than 65 books and countless articles, ranging over the fields of science, mathematics, philosophy, literature, and conjuring.

[3] The puzzle is the "river crossing problem" (we will return to this puzzle in chapter 12 of this book): *A man has to take a wolf, a goat, and some cabbage across a river. His rowboat has enough room for the man plus either the wolf or the goat or the cabbage. If he takes the cabbage with him, the wolf will eat the goat. If he takes the wolf, the goat will eat the cabbage. Only when the man is present are the goat and the cabbage safe from their enemies. How should the man carry the wolf, goat, and cabbage across the river?*

can only be learned by solving problems; however, this activity must be supported by strategies provided by an instructor. These general strategies would allow for solving new, yet unknown, problems in the future.

2. *Simplicity*: Educational puzzles should be easy to state and easy to remember. This is also very important, as easy-to-remember puzzles increase the chance that the solution method (which includes some universal mathematical problem-solving principles) is also remembered.
3. *Eureka factor*: Educational puzzles should frustrate the problem-solver! A puzzle should be interesting because the result is counter-intuitive: problem-solvers usually use intuition to start their quest for the solution and this approach usually leads them astray … Eventually a Eureka! moment is reached (Martin Gardner's Aha!) when the correct path to solving the puzzle is recognized. The Eureka moment is accompanied by a sense of relief, the frustration that was felt during the process dissipates, and the problem-solver may feel a sense of reward at their cleverness for eventually solving the puzzle. The Eureka factor also implies that educational puzzles should have elementary solutions that are not obvious.
4. *Entertainment factor*: Educational puzzles should be entertaining; otherwise it is easy to lose interest in them! Entertainment is often a side-effect of simplicity, frustration, the Eureka factor, and an "interesting" setting (e.g., the casino environment, a fight against dragons, dropping eggs from a tower).

Of course, we do not need to satisfy all of these criteria. For example, the zebra puzzle (puzzle 5.4), – not to mention the monkey and the rope puzzle (puzzle 12.28) – are impossible to remember as they contain too many details. Some puzzles (e.g., puzzle 6.2 on the traveling salesman problem) have no entertainment value, but there is no question they are educational! And a few puzzles, such as the 7-Eleven problem (puzzle 12.6) or some versions of *Nim* games (puzzle 11.6), do not have elementary solutions. Thus, in this book we have focused on educational puzzles using our own intuition and many years of teaching experience.

Besides being a lot of fun, the puzzle-based learning approach does a remarkable job of convincing students that (a) science is useful and interesting, (b) the basic courses they are taking are relevant, (c) mathematics is not *that* scary (there is no need to hate it!), and (d) it is worthwhile to stay in school, get a degree, and move into the real world which is loaded with interesting problems (problems perceived as real-world puzzles). These points are important, as most students are unclear about the significance of the topics covered during their studies. Oftentimes, they do not see a connection between the topics taught (e.g., linear algebra) and real-world problems, and they lose interest with predictable outcomes.

There are other well-established learning methodologies that address some of the above issues; these include problem-based learning and project-based learning (e.g., Blumenfeld et al. 1991, Bransford et al. 1986). Note, however, that the problem- and project-based approaches deal with quite complex situations where there is usually no single clear, unique, or correct way of proceeding. For example, projects may include assignments such as: *Where is the best location for a new airport in our city?* Or: *How can we run an efficient marketing campaign for a new product with a limited budget?* There may not be a single "best" solution to these problems or projects.

The emphasis in these approaches is usually on how to deal with the complexity of the problem and how to integrate the use of a wide range of techniques. Furthermore, project-based learning may involve teams of people with perhaps different specialist knowledge. With both problem- and project-based learning, a major piece of work is conducted under the supervision of an experienced facilitator acting in a mentoring role.

In a complementary contrast to problem-based learning, puzzles tend to be at the other end of the spectrum. They appear to be deceptively simple and usually have a single correct answer. An important part of completing a puzzle is to understand what we have learned by solving the puzzle and how we can apply this knowledge to other problems.

This book is the result of many years of experience in educating young engineers, mathematicians, computer scientists, and businessmen at many universities in many countries (USA, Mexico, Argentina, New Zealand, Australia, South Korea, Japan, China, Poland, Sweden, Germany, Spain, Italy, France, UK). Limited experiments using puzzle-based learning with these students have already produced outstanding course evaluations and countless comments that praise the problem-solving orientation of the course. We believe that the main reasons behind most students' enthusiasm for puzzle-based learning are:

- Puzzles are educational, but they illustrate useful (and powerful) problem-solving rules in a very *entertaining* way.
- Puzzles are engaging and thought-provoking.
- Contrary to many textbook problems, puzzles are not attached to any chapter (as is the case with real-world problems).
- It is possible to talk about different techniques (e.g., simulation, optimization), disciplines (e.g., probability, statistics), or application areas (e.g., scheduling, finance) and illustrate their significance by discussing a few simple puzzles. At the same time, the students are aware that many conclusions are applicable to the broader context of solving real-world problems.

We have organized this book in the following way: We begin with the *Introduction* (what a nice section to start with!), which explains in more detail the motivation behind this text. This is followed by thirteen chapters that are grouped into three parts. Part I presents the first three chapters, each of which discusses a simple problem-solving rule. Needless to say, each rule is illustrated by a collection of the best puzzles we could find. Part II presents eight chapters, from 4 to 11. These chapters cover various aspects of problems and problem solving by discussing constraints, optimization, probability, statistics, simulations, pattern recognition, and strategy. This part of the text makes a clear connection between various puzzles and different branches of science. It also includes a discussion on many mathematical problem-solving principles. Part III, on the other hand, consists of just two chapters that can be used as assignments (a collection of puzzles with and without solutions, respectively). These chapters include many puzzles that illustrate the applicability of various problem-solving rules and mathematical principles in a variety of domains.

Lastly, and most importantly, we would like to thank everyone who made this book possible, and who took the time to share their thoughts and comments on the subject of problem solving.

In particular, we would like to express our gratitude to a few individuals from the University of Adelaide: the Pro Vice Chancellor for Research Strategy Mike Brooks, the Executive Dean of the Faculty of Engineering, Computer and Mathematical Sciences Peter Dowd, the Associate Dean for Learning and Teaching of the Faculty of Engineering, Computer and Mathematical Sciences Mark Jaksa, and the Head of School of Computer Science Dave Munro for their encouragement and support during the execution of this whole project. Several faculties from different schools of the University of Adelaide helped us in this project. We thank Matthew Roughan, Nigel Bean, David Butler, Gary Glonek, and David Green from the School of Applied Mathematics, Ralf Zurbrugg from the School of Finance, Derek Abbott from the School of Electrical Engineering, Brad Alexander, Nick Falkner, Charles Lakos from the School of Computer Science for their comments on this text.

Thanks are also due to Alan Noble from Google, David Lindley from the Australian Computer Society, Peter Tischer from Monash University, Geoff Robinson from CSIRO, Anthony Harradine from Noel Baker Center for School Mathematics at Prince Alfred College in Adelaide, Ed Meyer from the Baldwin-Wallace College (Ohio, USA), John Woodward from the University of Nottingham, Stuart Brock from the Victoria University of Wellington, Chris Handley from the University of Otago, Jacek Koronacki and Antoni Mazurkiewicz from the Institute of Computer Science Polish Academy of Sciences, and Raja Sooriamurthi from Carnegie Mellon University for their comments, suggestions, and insights.

In most cases it is very difficult to trace the origin of a puzzle and give full credit to the inventor. Many puzzles (often in slightly different form) have surfaced many times in many different places, while others were simply passed on as word of mouth. This notwithstanding, we would like to acknowledge several puzzles that were published earlier in a variety of sources; these include one of the author's earlier books, *How to Solve It: Modern Heuristics* (some of the puzzles included in this text were "discovered" by the first author when he was nine years old, and then used many years later to "torture" the second author …). Many puzzles were found in journals (e.g., *The American Mathematical Monthly* or *Scientific American*), while others were adapted from books by Martin Gardner, *My Best Mathematical and Logic Puzzles* and *Entertaining Mathematical Puzzles*, and from other books: *How to Lie with Statistics*, by Darrell Huff; *Which Way Did the Bicycle Go?*, by Joseph D. E. Konhauser, Dan Velleman, and Stan Wagon; *Fifty Challenging Problems in Probability with Solutions*, by Frederick Mosteller; *Mathematical Puzzles: A Connoisseur's Collection*, by Peter Winkler; *The Moscow Puzzles*, by Boris A. Kordemsky; *Puzzles for Pleasure*, by Barry R. Clarke; *Innumeracy: Mathematical Illiteracy and Its Consequences*, by John Allen Paulos; *One Hundred Problems in Elementary Mathematics*, by Hugo Steinhaus; *The Lady or the Tiger? and Other Logic Puzzles* by Raymond Smullyan. Some puzzles were found in books only published in Poland and Russia (see the references at the end of this book), and to our best knowledge, there were no English translations of these works.

We would also like to thank the most famous fictional detective of all time, Sherlock Holmes, for providing us with the entertaining quotes at the beginning of each chapter. Mr. Holmes remains one of the most famous problem solvers of all time and his methodology is based on many interesting problem-solving rules: "*It is a capital mistake to theorize before you have all the evidence*"; "*When you follow two separate chains of thought, Watson, you will find some point of intersection which should approximate the truth*"; "*When you have eliminated the impossible, whatever remains,* *however improbable, must be the truth*"; and "*Singularity is almost invariably a clue. The more featureless and commonplace a crime is, the more difficult it is to bring it home*." Needless to say, his methodology bears a striking resemblance to the rules and principles presented in this text. Enjoy!

Adelaide, Australia
May 2008

Zbigniew Michalewicz
Matthew Michalewicz

Contents

Introduction

"Come, Watson, come!," he cried. "The game is afoot.
Not a word! Into your clothes and come!"

The Adventure of the Abbey Grange

How to solve it? This question is the holy grail of many disciplines – from mathematics and engineering, through to the sciences and business. We are constantly faced with this question during our lifetimes, both in the work environment and at home. *How much money should we invest? What are the best connections when flying from Australia to Europe? How should we schedule operations in the factory to minimize cost, while satisfying due dates and other requirements?* All these represent "problems" which require some solutions … hence the question: *How to solve it?*

Over the years, two primary approaches to problem solving have emerged. One is the *technical* approach (represented in many textbooks), which concentrates on specific problem-solving techniques. The other is the *psychological* approach, which is based on structural thinking – meaning that some structure is imposed on the thinking process during the problem-solving activity.

Let us discuss these two approaches in a bit more detail; for that purpose we have selected two popular texts. The first one is *Operations Research: An Introduction* by Hamdy A. Taha, and the other is a book by Edward de Bono, *Six Thinking Hats*. The first book illustrates the technical approach very well, as it is loaded with mathematical techniques for a variety of different problems. On the other hand, the second book presents a particular way of thinking. Let us have a closer look at these two books.

Operations Research: An Introduction by Hamdy A. Taha consists of several chapters, each of which relate to a specific problem type. For example, there is a chapter on linear programming, which is a particular technique for solving problems with many variables and where the objective and the values of these variables are expressed as linear expressions (puzzle 3.1 provides an example of a problem well suited for the linear programming approach, which is later outlined in chapter 6). Another chapter of Taha's book discusses a transportation model and its variants, while another presents a series of techniques applicable to network models (you should not be discouraged by this technical terminology – we only use it to make a point). There are chapters on goal programming, integer linear programming, dynamic programming, inventory models, forecasting models, etc. Each chapter includes selected references and a problem set.

For example, the chapter on inventory models includes the following exercise:

> McBurger orders ground meat at the start of each week to cover the week's demand of 300 lb. The fixed cost per order is \$20. It costs about \$0.03 per lb per day to refrigerate and store the meat. (a) Determine the inventory cost per week of the present ordering policy. (b) Determine the optimal inventory policy that McBurger should use, assuming zero lead time between the placement and receipt of an order. (c) Determine the difference in the cost per week between McBurger's current and optimal ordering policy.

Clearly, the problem is well-defined and very specific. Earlier parts of the chapter on inventory models discussed a general inventory model (where the total inventory cost is given as a total of purchasing cost, setup cost, holding cost, and shortage cost) and the classic economic order quantity models. The formula is derived in the chapter to provide the optimum value of the order quantity y (number of units) as a function of setup cost K associated with the placement of an order (in dollars per order), demand rate D (in units per time unit), and holding cost h (in dollars per inventory unit per time unit). The model suggests to order:

$$y = \sqrt{2KD/h}$$

units every y/D time units. Again, it is not our goal to scare you by providing a formula in the introductory part of this text (especially that the derivation of this formula requires some calculus), but rather to point out the specific nature of the problem and the specific (and very precise) solution. This example is a perfect illustration of the technical approach.

It seems that Taha's text is similar to many other texts from disciplines such as engineering, mathematics, finance, and business, in that it has two main characteristics:

(a) the problem types and corresponding techniques are very specific; and

(b) mathematics is used extensively.

However, there is usually no discussion on "how to solve a problem" – the text gives some formulas on how to arrive at a solution once the problem has already been reduced to the problem type defined in the text. As mentioned in the Preface, students are constrained to concentrate on textbook questions at the back of each chapter, using the information learned in that chapter.

There is nothing wrong with such texts – indeed, they are very useful in the classroom environment and make good textbooks for a variety of different courses. After all, students should master the appropriate techniques/methods/algorithms/etc. as this is expected from the educational system. In other words, the students are taught *how* to apply particular methods to particular problems, but only within the context of knowing that these methods are appropriate for these particular problems. They almost never learn how to *think* about solving problems in general. The same observation applies to all levels of education: in elementary school children are taught how to multiply two numbers, as this is considered (and rightly so) one of the basic skills needed for further advancement. On the other hand, children are not taught *when* to multiply two numbers. So in many elementary texts you can expect to find problems of the type:

> It takes 48 hours for a rocket to travel from the Earth to the Moon. How long will this trip take if a new rocket is twice as fast?

whereas problems like:

> It takes 48 hours for a rocket to travel from the Earth to the Moon. How long will this trip take for two rockets?

which force a child to *think* (whether to multiply or divide 48 by 2, or whether it would still take 48 hours), are not included. So all these specialized texts (whether on probability, statistics, simulations, etc.) that represent the technical approach for problem solving, do not present a problem-solving methodology. They just provide (very useful) information on particular techniques for particular classes of problems.

So let us now turn our attention to the other book, Edward de Bono's *Six Thinking Hats*, which represents the psychological approach. As we have indicated earlier, the book suggests some structure for the thinking process during the problem-solving activity. In particular, each of six hats represents a particular function of the thinking process:

White Hat: collection of objective facts and figures
Red Hat: presentation of emotional view
Black Hat: discussion of weaknesses in an idea
Yellow Hat: discussion on benefits of the idea
Green Hat: generation of new ideas
Blue Hat: imposition of control of the whole process

The general idea is that instead of thinking simultaneously along many directions, a thinker should do one thing at the time. Edward de Bono explains it very clearly:

> The main difficulty of thinking is confusion. We try to do too much at once. Emotions, information, logic, hope and creativity all crowd in on us. It is like juggling with too many balls.
>
> What I am putting forward in this book is a very simple concept which allows a thinker to do one thing at a time. He or she becomes able to separate emotion from logic, creativity from information, and so on. The concept is that of the six thinking hats. Putting on any one of these hats defines a certain type of thinking.

It seems that *Six Thinking Hats* is characterized by two facts (as are many other texts on thinking processes, which includes texts on critical thinking, constructive thinking, creative thinking, parallel thinking, vertical thinking, lateral thinking, confrontational and adversarial thinking, to name a few):

(a) the problem types and corresponding "techniques" are *not* very specific. The approach is very general and it applies to most problems (as opposed to specific problem types); and

(b) the approach is mathematics-free.

Indeed, the examples given in *Six Thinking Hats* vary from house selling activities, to advertising and marketing issues, to pricing products. Furthermore, mathematics is non-existent despite the fact that some problems may require more precise mathematics. There is no question that the approach proposed by Edward de Bono is very useful and that many corporations benefited from the *Six Thinking Hats*. On the other hand, the rejection of mathematics in *Six Thinking Hats* expresses itself even in the author's statements, such as:

> In a simple experiment with three hundred senior public servants, the introduction of the Six Hats method increased thinking productivity by 493 percent.

Well, this is very impressive, but any person with any "critical thinking" skills (or some fancy for precision) may ask for clarifications:

- What is the definition of productivity (especially in cases of senior public servants)?
- How is productivity measured?
- How is an improvement in productivity measured (with such great precision)?

Indeed, these are very important questions, and we will discuss the issue of understanding all terms and expressions present in the description of a problem in the first chapter of this book (as this is a key issue and the starting point of all problem-solving activities). In the case of the public servants, did three hundred employees fill out forms that evaluated their (increased) productivity? If so, then this can be compared to an example provided by Darrell Huff in his book *How to Lie with Statistics*. The *San Francisco Chronicle* published an article entitled "British He's Bathe More Than She's" and the story supported the title with the following facts (based on a survey that asked people to report their hot-water usage, carried out over 6,000 representative British homes):

> The British male over 5 years of age soaks himself in a hot tub on an average of 1.7 times a week in the winter and 2.1 times in the summer. British women average 1.5 baths a week in the winter and 2.0 in the summer.

Darrell Huff, discussing this case, made an excellent (and very important) observation. He wrote:

> … the major weakness is that the subject has been changed. What the Ministry really found out is how often these people said they bathed, not how often they did so. When a subject is as intimate as this one is, with the British bath-taking tradition involved, saying and doing may not be the same after all.

It seems that the same argument can be applied to the public servants. Most likely, their productivity was measured in hours (i.e., the shorter the time to make a decision, the better). Edward de Bono explains:

> A major corporation used to spend twenty days on their multinational project team discussion. Using the parallel thinking of the Six Hats method, the discussions can now take as little as two days.

However, if this was the case, then it seems there is something fundamentally very wrong with the whole picture, as *the quality* of the decisions reached is completely ignored and not measured! We acknowledge that the time to arrive at solution is important (as time is money), but in many cases *the quality* of solution is the most important aspect.

There is an excellent book (on science and education, one can say) by Eliyahu M. Goldratt and Jeff Cox, *The Goal*. The book describes the struggle of a plant manager who tries to improve factory performance. He worries about productivity, excess inventories, throughput, balancing capacities, and many other measurements. Only with a help of a consultant does he realize that there is only one goal and one measurement:

> "The goal of a manufacturing organization is to make money and everything else we do is means to achieve the goal."

Similarly, in the problem-solving processes there is only one goal: to find the best possible solution. Of course, very often there is a tradeoff between the time needed to find a solution and the quality of the solution (this is often discussed in computer science courses on analysis of algorithms), but it seems that the *Six Thinking Hats* method is concerned with only the secondary aspect of problem-solving: time efficiency. Precise evaluation of the solution is of lesser importance.

Thus the psychological approach looks like the *opposite extreme* of the technical approach in the spectrum of problem-solving methodologies, as the former focuses on organizational issues of "thinking" for general problems, rather than specific techniques on how to arrive at a solution. Furthermore, the psychological approach uses natural language to describe its mechanisms, whereas the technical approach uses mathematics as a problem-solving language.

Which of these two approaches (technical versus psychological) should be used in the real world? Well, each of these two approaches has a crowd of enthusiasts and supporters; however, it seems that the technical approach is based on the solid fundamentals of science. Even some philosophers and psychologists tend to agree. One of the pearls of wisdom taught by Anthony de Mello in his famous book, *One Minute Wisdom*, was the following observation:

> Better to have the money than to calculate it; better to have the experience than to define it

It is easy to extend the above statements (while preserving their spirit) by stating that:

> Better to have the problem-solving skills than to discuss them.

On the other hand, representatives of the technical approach admit that:

> Although mathematics is a cornerstone of Operations Research, one should not 'jump' into using mathematical models until simpler approaches have been explored. In some cases, one may encounter a 'commonsense' solution through simple observations. Indeed, since the human element invariably affects most decision problems, a study of the psychology of people may be key to solving the problem. (Hamdy A. Taha, *Operations Research: An Introduction*)

These comments are followed by a delightful example, where the problem of slow elevator service in a large office building was solved not with mathematical queuing analysis or simulation, but by installing full-length mirrors at the entrance to the elevators: the complaints disappeared as people were kept occupied watching themselves (and others) while waiting for the elevator!

There are many merits in concepts related to critical, vertical, lateral, and other thinking paradigms. We will see in the following chapters in this text that the ability to ask the right (critical) questions, the ability to follow a (vertical) line of thought, and the ability to think laterally (out of the box) are essential in the process of problem solving. However, mathematics – the queen of all sciences – must remain the universal language of problem solvers. Otherwise, as we saw, there is a danger of making imprecise statements, and what is worse, there is a danger of finding (and implementing) poor solutions! In this text we have tried to combine these two approaches: despite the fact that the text is elementary, we have used mathematical notation (as simple as possible) all the way through. At the same time, we have introduced a few problem-solving rules (that are related to various categories of thinking) to guide the process.

Interestingly, puzzle-based learning mixes different learning paradigms together. Twenty-five centuries ago Confucius[4] said:

> By three methods we may learn wisdom: first, by reflection, which is noblest; second, by imitation, which is easiest; and third, by experience, which is bitterest.

Indeed, puzzle-based learning allows us to learn problem-solving skills by *all* the above methods. We learn by experience (as we can learn problem-solving skills only by solving problems). We learn by imitation, as it is helpful to imitate (apply) some principles and techniques. And above all, we learn by reflection, as puzzle-based learning encourages us to reflect on:

- What are we learning?
- How are we learning it?
- How are we using what we have learned?

There are also other approaches proposed in the past that address the key question: "How can I get my students to think and solve problems?" The problem-based learning approach proposed in the 1960s at McMaster University Medical School (Hamilton, Ontario, Canada) is an instructional method that challenges students to "learn to learn," working cooperatively in groups to seek solutions to real-world problems. Problem-based learning is aimed at enhancing content knowledge and fostering the development of communication, problem-solving, and self-directed learning skills. It has since been implemented in various undergraduate and graduate programs around the world.

Today the defining characteristics of problem-based learning are:

- Learning is driven by challenging, open-ended problems.
- Students work in small collaborative groups.
- Teachers take on the role of "facilitators" of learning.

Accordingly, students are encouraged to take responsibility for their group and organize and direct the learning process with support from a tutor or instructor. In other words, problem-based learning is any learning environment in which the problem drives the learning. That is, before students learn some knowledge they are given a problem. The problem is posed so that the students discover that they need to learn some new knowledge before they can solve the problem. Student participation involves hands-on investigative/laboratory activities that develop inquiry and intellectual skills. These activities give students an opportunity to appreciate the spirit of science and promote the understanding of the nature of learning.

A classic example of problem-based learning is the famous "Egg-Drop" experiment which has

---

[4] A Chinese thinker and social philosopher (551 BC – 479 BC), whose teachings have influenced thought and life of millions of people of Far East.

been a standard in science instruction for many years. In this experiment students are asked to construct some type of container that will keep a raw egg from cracking when dropped from ever-increasing elevations. A number of different groups can be set up to search for ways of approaching this problem. Students will be confronted with some long-standing and resilient misconceptions concerning free-fall (for instance, that heavy objects fall to the earth quicker/slower than lighter objects). By encouraging experimentation and communication of their results, some students may see the need to use mathematics in their approach to this problem – however, many students would stay with intuitive solutions.

Students may come to value the notion of a prototype as they take part in the design process, and their investment in the project should increase accordingly. The solution presented for this project can be either a group or individual accomplishment depending on how the instructor wishes the dynamics of the class to develop.

But puzzle-based learning offers a very different intellectual feast for the "Egg-Drop" experiment. Suppose you wish to know which floors in a high building are safe to drop eggs from in a special container and which floors will cause the eggs to break upon landing? We can eliminate chance and possible differences between different eggs (e.g., one egg breaks when dropped from the 7th floor and another egg survives a drop from the 20th floor) by making a few (reasonable!) assumptions:

- An egg that survives a drop can be used again (no harm is done and the egg is not weaker).
- A broken egg cannot be used again for any experiment.
- The effect of a fall is the same for all eggs.
- If an egg breaks when dropped from some floor, it would break also if dropped from a higher floor.
- If an egg survives a fall when dropped from some floor, it would survive also if dropped from a lower floor.

Obviously, if only one egg is available for experimentation to determine the first egg-breaking floor, we have to start with dropping the egg from the first floor. If it breaks, we know the answer. If it survives, we drop it from the second floor. And we continue upward until the egg breaks. Clearly, the worst-case scenario would require as many drops as the number of floors in the building. Now, the challenge begins when we have two available eggs. What is the least number of egg drops required to determine the egg-breaking floor?

To solve this problem, no laboratory is required: just basic problem-solving skills plus the ability to add and subtract numbers! We believe that this puzzle-based version of the "Egg-Drop" problem is of equal intellectual value and complements the original "Egg-Drop" experiment offered by the problem-based learning approach.[5]

Since problem-based learning starts with a problem to be solved, students working in a problem-based learning environment should be skilled in problem solving or critical thinking or "thinking on your feet" (as opposed to rote recall). Many educators believe that some qualifying examinations – in which the problem solving (thinking skills) of the candidates are tested – should be conducted before the candidates are admitted. In the McMaster University Medical School, one of the five

---

[5] This problem is discussed fully in chapter 6 of this text (puzzle 6.7).

criteria for admission is a test of the candidates' problem-solving skills. Unfortunately, many universities introduce problem-based learning courses without pre-screening or developing their students' skills in problem solving. So a puzzle-based learning course (or unit) fits very well as a prerequisite for later problem-based learning activities.

As stated in the Preface, the lack of problem-solving skills in general is the consequence of decreasing levels of mathematical sophistication. People (again, in general) have difficulties dealing with numbers, to say nothing of basic mathematical concepts! There is a great book written by John Allen Paulos, *Innumeracy: Mathematical Illiteracy and Its Consequences*, where the author demonstrates how much mathematical ignorance pervades both our private and public lives and results in misinformed government policies, confused personal decisions, and an increased susceptibility to pseudo-sciences of all kinds. The book is largely concerned, in the author's words, with "… a lack of numerical perspective, an exaggerated appreciation for meaningless coincidence, a credulous acceptance of pseudo-sciences, an inability to recognize social trade-offs, and so on …"

Indeed, it is a scary picture when a university student argues that hair does not grow in miles per hour or an educated grown-up believes that if there is a 50 percent chance of rain on Saturday and 50 percent chance of rain on Sunday, then there is a 100 percent chance of rain during the weekend (these examples were taken from John Paulos's book). A version of the latter example was turned into a joke, where a travel agent advises a male traveler to date only women with brown hair while in a particular country, as statistics about women in that country are very clear: 50 percent of the women have brown hair and 50 percent of women suffer from tuberculosis! Such mathematical ignorance may explain a growing popularity of psychological approaches for problem solving, but this does not seem the right way to address problems …

To make sure this text is not beyond the understanding of readers who are not well versed in mathematics, we have assumed an elementary level of mathematical skills. In fact, basic knowledge of high school mathematics is more than sufficient to follow the whole text. We also believe that mathematical notation used in this text will not spoil the enjoyment of solving many entertaining puzzles! Further, we tried to convince the reader that mathematics is *not* just a bunch of techniques invented in 19th century and before. New mathematics is constantly being generated – but it is impossible to teach how to generate "new" mathematics. It comes down to solving puzzles and inventing new techniques to do so.

Let us conclude this introduction with the following observation. Numerous mathematicians have put a lot of effort into finding a middle ground between the technical and psychological approaches to problem solving. The best known work, without a doubt, is Gyorgy Polya's *How to Solve It*, which stands out as one of the most important contributions to problem-solving literature of the 20th century. Even after moving into the new millennium the book continues to be a favorite among teachers and students for its instructive methods. Other works include *I Hate Mathematics* written by Marilyn Burns, which is full of tips and methods for solving problems.

Another trend represented by several mathematicians is based on the belief that puzzles (usually mathematical puzzles) are quite educational and that we should educate students by incorporating puzzles into various curricula. Probably the unquestioned leader of this trend is Martin Gardner, who collected and published thousands of fantastic puzzles – on all levels – in his books (e.g., *My*

Best Mathematical and Logic Puzzles, *Entertaining Mathematical Puzzles*, *The Colossal Book of Mathematics*, or *The Colossal Book of Short Puzzles and Problems*) and various journals (e.g., he ran a puzzle column in *Scientific American* for many years).

Many other mathematicians were also believers in this approach. Joseph Konhauser, while at Macalester College, published *Problem of the Week* for 25 years to attract students' interest as his problems (or rather puzzles) which had special appeal and often some surprising twists. His best puzzles were published in the volume *Which Way Did the Bicycle Go?* by Joseph D. E. Konhauser, Dan Velleman, and Stan Wagon. The famous Polish mathematician, Hugo Steinhaus, published a collection of entertaining puzzles in the volume *One Hundred Problems in Elementary Mathematics*; the American mathematician Frederick Mosteller wrote *Fifty Challenging Problems in Probability with Solutions*, and the German mathematician Arthur Engel published *Problem-Solving Strategies*, a volume that includes over 1,300 examples and problems. Peter Winkler also wrote *Mathematical Puzzles: A Connoisseur's Collection*, Boris A. Kordemsky published *The Moscow Puzzles*, and Barry Clarke: *Puzzles for Pleasure*. And the list goes on.

We wholeheartedly support this trend and direction, and believe this book provides an important contribution. And with all these remarks, clarifications, and explanations, we are ready to proceed.

PART I

Rules 1—2—3

1 The Problem: What are you after?

"I confess that I can make neither head nor tail of it. Don't you think that you have kept up your mystery long enough, Mr. Holmes?"

Silver Blaze

To illustrate one of the main points of this chapter, let us introduce the first puzzle:

Puzzle 1.1 At six o'clock in the morning the wall clock struck 6 times and the time between the first and last strokes was 30 seconds. How long will the clock take to strike 12 times at noon?

It seems that one can answer this puzzle without much thought. If 6 strikes take 30 seconds, then a single strike "takes" 5 seconds. Hence 12 strikes will take 12 × 5 = 60 seconds, a full minute. Right? Well, not exactly. Problem-solving activities require some reasoning skills (and this is what this book is about), and the first skill required is the ability to understand the problem (which is the theme of this chapter).

Actually, the problem is not that easy: as a matter of fact, if we do not make any additional assumptions, there would be no unique solution! Simple (but careful) reasoning should convince us that this is the case. To start with, note that between any two consecutive strikes there is a break. So if the clock struck 6 times, then the time between the first and last strokes in the puzzle is really the total of: (a) the time for all 6 strokes, and (b) the time for 5 breaks in between strokes. If x and y represent the times required for a single stroke and for a break in between two strokes, respectively, then the information given in this puzzle can be written as:

$$6x + 5y = 30$$

The question is, on the other hand, how long will the clock take to strike 12 times? As there would be 11 breaks in the sequence of 12 strokes, the equation is:

$$12x + 11y = ?$$

Of course, the first equation leads us to the following:

$$12x + 10y = 60$$

so we know that the time between the first and last strokes would take more than 60 seconds (by one y, which is the time for one break), as:

$$12x + 11y = (12x + 10y) + y = 60 + y$$

If we do not know the length of a break, it would be impossible to solve the problem. For example, if it takes 1 second to strike (i.e., $x = 1$) then a break takes 4.8 seconds (i.e., $y = 4.8$) as $6x + 5y = 30$. In this case, it will take 64.8 seconds for 12 strikes. If, on the other hand, it takes 2 seconds to strike (i.e., $x = 2$), then a break takes 3.6 seconds (i.e., $y = 3.6$). In this case it will take 63.6 seconds for 12 strikes. As you can see, the puzzle has *many* possible solutions …

To get a unique solution we have to assume something. We may, for instance, assume that the strikes take no time (i.e., $x = 0$). Then 5 breaks take 30 seconds, so the break between two strikes takes 6 seconds. For 12 strikes, there are 11 breaks between strikes, so it would take 66 seconds for the clock to strike 12 times.

In this puzzle it was essential to distinguish between strikes and breaks, and understand that it is not the case that a strike takes 5 seconds (in the case of $x = 0$) but rather that a break between two strikes takes 6 seconds. Without this understanding we may jump to the wrong conclusion.

The following puzzle also illustrates the process of understanding the problem:

Puzzle 1.2 You drive a car from point **A** to point **B** at a constant speed of 40 km/h, and, on arrival at point **B**, you immediately return to point **A**, but at a higher speed of 60 km/h. What was your average speed for the whole trip?

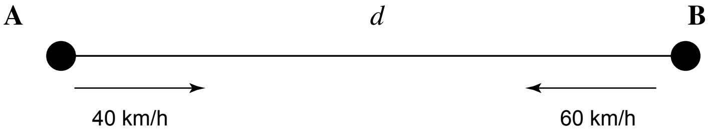

The vast majority of people say that the average speed for the whole trip was 50 km/h, as they do not have a clear understanding of the term “average.” Furthermore, the correct answer of 48 km/h seems very counterintuitive to them!

How it is possible that 48 km/h is the correct answer? After all, the average of 40 and 60 *is* 50. To see it clearly, we have to understand the question (i.e., the problem we are trying to solve). Note that the term “average speed” is defined as a ratio between the distance and time, thus:

$$v_{avg} = D/T$$

where D and T represent the *total* distance and *total* time of the whole trip, respectively. There are no other definitions of “average speed” for a trip! In all circumstances, while calculating average speed, we need to know the total distance traveled and the total time for the whole trip. Other reasoning, like averaging numbers 40 km/h and 60 km/h, are simply wrong!

In our case, the total distance D is equal to $2d$, where d is the distance between points **A** and **B** (as the trip takes us from point **A** to point **B** and back). On the other hand, the total time T for the whole trip consists of two components: the time t_{AB} to get from point **A** to point **B** and the time t_{BA} to get from point **B** to point **A**. Thus T is equal to $t_{AB} + t_{BA}$. So, the average speed for the whole trip is given as a ratio:

$$v_{avg} = D/T = 2d/(t_{AB} + t_{BA})$$

How long would it take us to drive from point **A** to point **B** if we drove at constant speed of 40 km/h? Note that the travel time is a ratio between the distance and the speed, so:

$$t_{AB} = d/40$$

Similarly:

$$t_{BA} = d/60$$

With this in mind, we easily arrive at the final answer:

$$v_{avg} = 2d/(t_{AB} + t_{BA}) = 2d/(d/40 + d/60) = 2/(1/40 + 1/60) = 48$$

The above two puzzles illustrate the most important point in all problem-solving activities: we should have a solid understanding of the problem before we attempt to solve it! This includes understanding the concepts that we are dealing with. The term "average" is especially confusing.[6] Consider another example that shows how important the definition of "average" is in many circumstances: There is a common perception that "we always join the slower line." If we imagine two lines of equal length and assume that in the absence of any other information they will behave (on average) the same, then the probability of joining the slower line should be 50 percent. So where does this intuition come from? The reason is the following: the 50 percent probability of joining the slower line is an *event-average*, whereas the intuition above is based on the *time-average*. For example, if we calculate the probability of *being* in the slower line (rather than just *joining* it) we will (on average) spend more time in the slower line, and this time-average will result in us having a higher probability of *being* in this line.

So the precise definition of "average" being used in a problem is really important and not always fixed in stone (as the above example illustrated). This discussion leads us to the first rule of problem solving, which we can formulate as:

> **Rule #1**. *Be sure you understand the problem, and all the basic terms and expressions used to define it.*

Note that quite often we use a variety of terms, like: "average," "middle," "larger," "better," without a proper understanding of these terms in the context of the problem at hand. Apart from understanding these terms, confusion can also arise from the way a problem (or a situation, event,

---

[6] The term "average" has a very loose meaning. For a beautiful discussion on how to take advantage of this imprecision, see chapter 2 of *How to Lie with Statistics* by Darrell Huff, who wrote: "It is a trick commonly used, sometimes in innocence but often in guilt, by fellows wishing to influence public opinion or sell advertising space. When you are told that something is an average you still don't know very much about it unless you can find out which of the common kinds of average it is – mean, median, or mode."

etc.) is described. You have probably heard people say "the summer of 2001 was much nicer than the summer of 2002" or "the students in my class this year were smarter than the students last year." We usually know (intuitively) what such statements mean, but upon closer inspection we might find our intuition giving way to some lingering doubts about how exactly we should interpret these sorts of claims. The following story illustrates **Rule #1** very well.

Two groups of students are attending school. The students in group **A** boast that they are taller than the students in group **B**, while the students in group **B** enjoy the reputation of being smarter than the students in group **A**.

One day, one of the students from group **A** approached a student from group **B**, and said "We are taller than you!" The student from group **B** thought about this statement and replied: "What do you mean that statement? Do you mean that:[7]

1. Each ***a*** is taller than each ***b***?
2. The tallest ***a*** is taller than the tallest ***b***?
3. Each ***a*** is taller than some ***b***?
4. Each ***b*** is smaller than some ***a***?
5. Each ***a*** has a corresponding ***b*** (and each of them a different one) whom he surpasses in height?
6. Each ***b*** has a corresponding ***a*** (and each of them a different one) by whom he is surpassed?
7. The shortest ***b*** is shorter than the shortest ***a***?
8. The shortest ***a*** exceeds more ***b***'s than the tallest ***b*** exceeds ***a***'s?
9. The sum of heights of the ***a***'s is greater than the sum of heights of the ***b***'s?
10. The average height of the ***a***'s is greater than the average height of the ***b***'s?
11. There are more ***a***'s who exceed some ***b*** than there are ***b***'s who exceed some *a*?
12. There are more ***a***'s with height greater than the average height of the ***b***'s than there are ***b***'s with height greater than the average height of the ***a***'s?
13. Or that the median height of the ***a***'s is greater than that of the ***b***'s?"

These are excellent questions, and the student from group **A** could only reply: "I'm not sure. I'll need to think about that …"

Indeed, there is something to think about. Not only can the expression "We are taller than you!" be interpreted in many different ways (and the above list of 13 different interpretations is far from complete, as we can invent more sophisticated – and sometimes crazy – interpretations, such as: the product of heights of the ***a***'s is greater than the product of heights of the ***b***'s), but these interpretations are not independent from each other, as some of them are stronger in the sense they imply other interpretations. For example, the answer "yes" to question #1 (i.e., each ***a*** is taller than each ***b***?) implies "yes" to question #2 (i.e., the tallest ***a*** is taller than the tallest ***b***?) – this is obvious. It might be less obvious to find all pairs of questions from the above list of 13 questions,

---

[7] For simplicity, ***a*** denotes a single student from group **A** and ***b*** denotes a single student from group **B**.

where an answer of "yes" to the first question implies the answer "yes" to the second question. If we find these pairs, we would also find questions that are equivalent in a sense that the answer to both questions must be the same.[8]

Rule #1 has some powerful consequences, as the process of understanding all the terms and expressions used in the description of the problem can often be the hardest part of the problem-solving activity! Especially given that real-world problems are usually described in imprecise terms. When people communicate with one another, they rarely resort to a high level of exactness. This simply is not practical. For the question "What time is it?" we should not expect the answer: "It is 5:43:27." People understand and operate on terms based on their own individual understanding of the degree to which those terms represent some particular condition. Words like "low," "high," "close," "very old," "red," "early," and so forth, each have a general meaning that we understand. Those meanings are imprecise, but useful nevertheless. Trying to impose a precise meaning to each term is not only impractical, it is truly impossible, because each term means something different to each person.

The classic example, used quite often by Lotfi Zadeh (the scientist who invented fuzzy logic) in his presentations, was to quote some article:

> Late afternoon was partially cloudy, yet there were many people on the streets of this large city.

In this statement, what does "*Late afternoon*" mean? Would 3:25 p.m. qualify as a late afternoon hour? What does "*partially cloudy*" mean? Are there three or four clouds on the sky? More? Less? It seems this is a bit unclear. But what does "*many people*" mean? Forty? Fifty? More? And what does "*large city"* mean? One million inhabitants? More? Less?

Problems are often expressed in natural (thus imprecise) language and so the process of understanding all the terms and expressions necessary for solving the problem is not always straightforward. A perfect example of this is given later (e.g., see puzzle 12.28); however, even the following simple puzzle illustrates the case quite well:

Puzzle 1.3

Mrs. Brown was celebrating her birthday, and one of the guests asked her about her age. Mrs. Brown replied that the total of her age and the age of her husband, Mr. Brown, is 140, and then she added: "My husband is twice the age I was when my husband was my age." How old is Mrs. Brown?

---

[8] If we denote the questions by the numbers they are designated with, and let the symbol p ® q denote the fact that the answer "yes" to question p implies the answer "yes" to question q, then the complete list of implications is:

1 ® 2, 1 ® 3, 1 ® 4, 1® 7, 1 ® 8, 1 ® 10, 1 ®11, 1 ® 12, 1 ® 13,
2 ® 4,
3 ® 7,
4 ® 2,
5 ® 7,
6 ® 2, 6 ® 4, 6 ® 9,
7 ® 3,
8 ® 3, 8 ® 7.

Thus, the answers to questions 2 and 4 are identical (as 2 ® 4 and 4 ® 2) and the same is true for the answers to 3 and 7.

The hard part of this puzzle is *not* finding the answer, but rather, trying to *understand* all the information given in the problem. It is not difficult to solve the problem once we understand it – once we know how to transform the information into the appropriate equations. We will return to this puzzle in chapter 3, where we discuss the importance of building a model of the problem.

Now let us turn our attention to a delightful puzzle about cats and rats. Contrary to the previous puzzle, the following one seems very clear; however, closer examination will reveal that not everything is as straightforward as it appears:

Puzzle 1.4 Three cats catch three rats in three minutes. How long it will take them to catch 100 rats?

Indeed, there is very little to think about here! The typical reasoning is as follows: If it takes three cats three minutes to catch three rats, then it takes them just one minute to catch one rat. So, these three cats catch one rat per minute, and would catch 100 rats in 100 minutes. Easy.

But take another look at the puzzle and the solution. We have made an important assumption without any justification (as this assumption was *not* stated in the problem): We have assumed that *all three cats* go after the *same rat*; it would take them one minute to catch it, and then they turn their attention to the next rat. But again, this was just our assumption. It might be that the cats' strategy is very different – it might be that each cat chases a *different* rat and it takes three minutes to accomplish the task. If this is the case, it would take them 6 minutes to catch 6 rats, 9 minutes to catch 9 rats, and so on. It would take them 99 minutes to catch 99 rats. But this was not the question: we have to find out, how long it would take them to catch 100 rats (and not just 99 of them)! Again, this is unclear and there is nothing in the problem description that can assist us in getting the answer. It might be that it takes them a full 3 minutes to catch the final rat (one cat goes after the rat and the other two cats just watch), so the total time would be 102 minutes. But maybe they can do it quicker, if all of them participate in the chase?

As we can see, to solve this puzzle we need to look deeper into the problem. What is the cats' strategy for catching rats? Do they work together or individually? We need this information to come up with the correct answer! Without this information, we can only guess. There is also another silent assumption in this puzzle about the linear relationship between the "number of units to complete" and the "amount of time needed for the operation." Note that it is quite common to assume that if one worker can accomplish a task in 12 hours, then 2 workers could complete it in 6 hours, and 3 workers in 4 hours. This might be true under some circumstances; however, only in very unusual circumstances could 12 workers complete the task in 1 hour, and it would be even harder to believe that 720 workers could complete the task in 1 minute. Obviously, there are some non-linear relationships between the number of workers and the total time to complete the task. Returning to our last puzzle, it might be quite easy for a cat to catch a rat (any rat) out of a crowd of 100 rats – this may require a short amount of time. In other words, if 3 cats can catch 3 rats in 3 minutes, it might be that the same 3 cats can catch 3 rats (but out of 100 rats) much faster: within one minute. Or it might take much longer, as a cat might chase a particular (selected) rat and the other rats are obstacles!

So even without using complex expressions (like in puzzle 1.3), a problem that contains only straightforward formulation might prove to be quite confusing and unclear after closer examination. We will return to puzzle 1.4 in chapter 9, where we look at some other possibilities for approaching this problem by simulating the environment. In the meantime, let us have a look at another puzzle, this time from the area of finance:

Puzzle 1.5 Mr. White sold his sailing boat to Mr. Brown for $10,000. A few weeks later, Mr. Brown discovered that the boat did not live up to his expectations and – being very disappointed – sold the boat back to Mr. White for $9,000. A few weeks later, Mr. White sold his boat again, this time to Mr. Green, for $9,500. The question is: What was Mr. White's total profit?

Before we try to solve this puzzle, let us apply **Rule #1**: *Be sure you understand the problem, and all the basic terms and expressions used to define it*. At first blush, there is nothing complex in the description: $10,000 is just ten thousand dollars; $9,000 is nine thousand … A sailing boat? Indeed, it might be any object of value (a car, a house), so we do not need to worry too much about the details of the sailing boat. What else? The puzzle includes the term "total profit." Usually, the term is defined as the difference between the cost of an object and its sale price. In other words, if Mr. White bought a sailing boat for $8,000 and later sold it for $10,000, the profit would be:

$10,000 – $8,000 = $2,000

With this realization, we should sense a problem: we do not know the original cost of the boat. The puzzle begins when Mr. White is already the owner of the sailing boat – there is no hint as to its original price (Mr. White might have inherited it). So, how should we define the term "total profit" when the original price is unknown? This is a perfect example of a problem that allows many reasonable answers, which depend on our interpretation of the term "total profit." Let us look at some of them:

- *Possibility #1*: The total profit is $500. Because the original price is unknown, we cannot determine the profit of the first transaction. However, we do know that the boat was later bought back for $9,000 and then resold for $9,500, so the total profit is $500.
- *Possibility #2*: The total profit is $1,000. The boat was sold for $10,000 and later bought back for $9,000. At this stage Mr. White is in the exact same spot as before (he is an owner of the same boat) except that he has extra $1,000 in his pocket. This represents his profit. Because we still do not know the original cost of the boat, we can not determine the profit (if any) of the second transaction. So the total profit is $1,000.
- *Possibility #3*: The total profit is $1,500. As before we can argue that the profit of Mr. White after the first transaction was $1,000. However, the second transaction generates a profit of $500, as Mr. White sold his boat for $500 more than he paid for it. So the total profit is $1,500.

Again, which of these is correct? All of them! Everything depends on our interpretation of the term "total profit." Probably the best answer is the following: "The total profit after all these transactions is \$10,500 minus the original cost of the boat." The reason is that Mr. White – after all these transactions – has \$10,500 in his pocket, but no longer has the sailing boat.

The next two puzzles illustrate a very interesting and important point: It might be that all the basic terms and expressions used in the description of the problem are clear, yet the description suggests a particular way of solving the problem – and of course, this silent suggestion is misleading.

Puzzle 1.6 There are six matches on the table. Your task is to construct four equilateral triangles where the length of each side is equal to the length of a match.

It is easy to construct two such triangles using five matches:

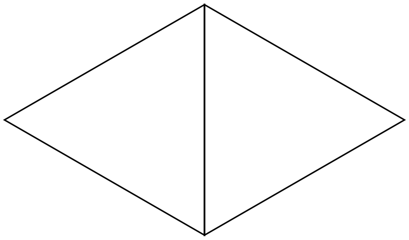

but it is difficult to extend this further into four triangles, especially as only one match remains. Here, the formulation of the problem is misleading because it suggests a two-dimensional environment (i.e., the matches were placed on a *table*). But to find a solution requires moving the problem into three dimensions:

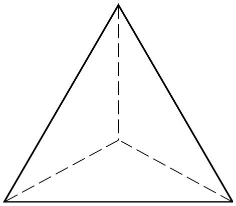

If we start with the wrong problem-solving environment, we will never find the right answer. So why do so many people – while solving this puzzle – find themselves in the wrong problem-solving environment? Because **Rule #1** is often overlooked! After all, the above problem is easy to understand, and all the basic terms and expressions (e.g., matches, table, and equilateral triangles) are quite basic … or are they? What about the phrase "to construct"? What does it mean, precisely? Does it mean that we have to position these matches on the surface of the table? If we are not sure, then we can easily enter the wrong problem-solving environment if we do not ask for clarifications!

One of the difficulties here was the (silent) suggestion that the construction phase should take place on the surface of the table – otherwise, there is no need to mention the table at all. The first sentence of the puzzle could have been: "You have six matches." So the term "table" serves as a false hint here. Note that if we give this problem to another person, we can experiment by giving them some additional hints like: you can break the matches into smaller pieces. Such hints are *harmful* in the sense they do not contribute anything toward obtaining the solution, and they expand the number of possibilities we have to consider.

The next puzzle is similar in the sense that it suggests some arrangements that may make the problem-solving process more difficult:

Puzzle 1.7 There are twelve matches; each of them is of the same length. You can arrange these matches in various polygon shapes, with each polygon having some area. For example, the three polygons constructed below:

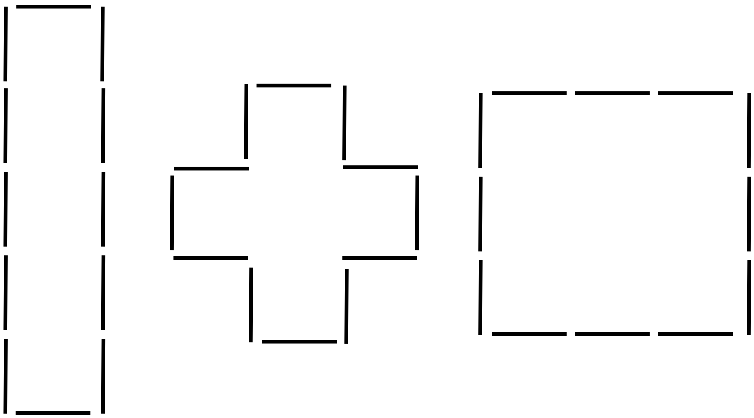

have areas of 5, 5, and 9 square units, respectively. Now, the problem is to use all 12 matches (and as in the examples given above, the entire length of each match must be used) to build a polygon with an area of 4.

After examining the examples given above, the solution is not so obvious. You may recall that the numbers 3, 4, and 5 are special in the sense that they represent lengths of sides of right-angled triangle:

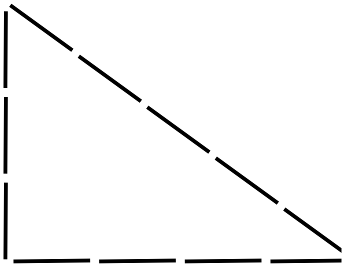

as $3^2 + 4^2 = 5^2$. The area of such triangle is 6. By making small modification (cutting off a small rectangle with the area of 2) we get a polygon with the required area of 4:

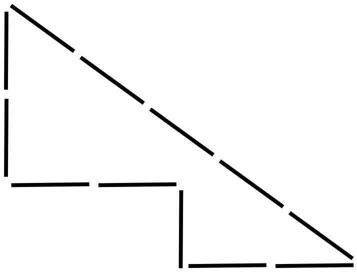

Again, **Rule #1** is useful here. The problem is that many people think the term "polygon" means "convex polygon"… hence the confusion[9]. For further discussion on this problem, see Martin Gardner's *My Best Mathematical and Logic Puzzles*.[10]

The final puzzle of this chapter is a classic, described in many books on problem-solving activities, thinking processes, etc. It also illustrates the concept of "thinking out of the box."

Puzzle 1.8 There are 9 dots placed on a surface:

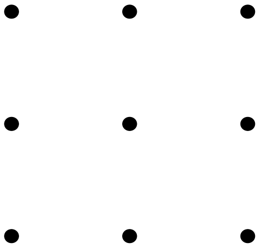

Connect all these dots with the smallest number of straight lines without lifting your pen from the paper.

Again, the puzzle just suggests that we should search for a solution within the four corner dots. If you do it, good luck! The best solution would consist of five straight lines:

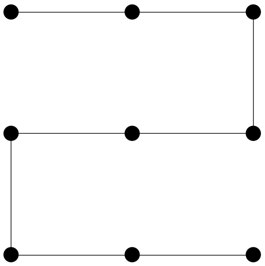

---

[9] A convex figure (polygon or not) is a one where by connecting two internal points by a straight-line leaves the whole straight-line segment inside the figure.

[10] Apparently, by constructing a star-shaped polygon and adjusting the width of the star's points we can produce any desired area between 0 and 11.196 – the latter number represents the largest possible area constructed from 12 matches: the area of a regular dodecagon …

To get a better solution, however, you have to literally "think of out the box":

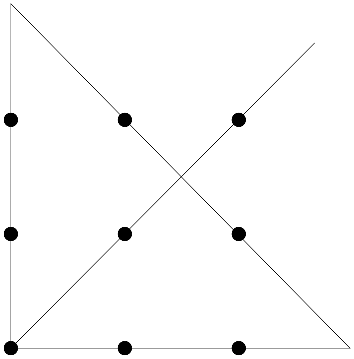

Only four line segments are required to connect all 9 dots. This puzzle illustrates **Rule #1** very well: do we really understand all the basic terms and expressions present in the description of the problem? Actually, the formulation of the problem is loaded with imprecision. The first sentence says: "There are 9 dots placed on a surface." So, what is a dot? Is it a round and relatively small object (as presented in the figure above) or rather, is this term used in the mathematical sense, where a dot is a point without any size? The figures above suggest the former interpretation, but in the process of solving this puzzle we somehow stick with the latter interpretation. So, if this is the case (i.e., each dot is indeed a small, round object), we can do even better: we can connect all 9 dots with just three line segments. But to do so, we have to think *really* "out of the box," as we have to extend these line segments far away from the "action square":

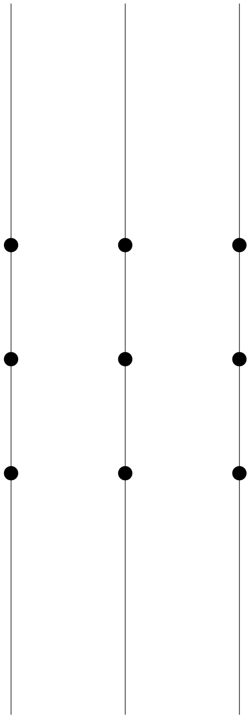

So, do we have our solution? Three line segments: that is not bad! But what about the other terms present in the description of the puzzle: "surface" and "line"? Do we mean "flat surface" or any surface? This is important as these 9 dots can be placed on a very special surface: a sphere … And if this is the case, they can be connected by a single line that takes the form of a spiral on the sphere. But in that case, can we call this spiral a line? Not in the mathematical sense …

Concluding remarks

All the puzzles discussed in this chapter illustrated a very simple, but important rule, which can be summarized as the follows:

> **Rule #1.** *Be sure you understand the problem, and all the basic terms and expressions used to define it.*

This rule stipulates that before attempting to solve any problem, spend as much time as possible understanding all the aspects of the problem. It pays off. This is similar to an old military proverb that states: "The more sweat during exercise, the less blood during battle." The same applies to problem-solving: the more time we spend on understanding (analyzing) the problem, the less time it takes to come up with the solution. Patience is a prerequisite for effective problem-solving.

The process of understanding the problem is also closely related to *critical thinking*, which, in a nutshell, refers to our ability to ask and answer critical questions related to what we have read or heard. A typical example is given in the introductory pages of *Asking The Right Questions: A Guide to Critical Thinking*, by M. Neil Browne and Stuart M. Keeley. Let us consider the following passage:

> Arguments for banning guns are mostly myths, and what we need now is not more laws, but more law enforcement. One myth is that most murderers are ordinary, law-abiding citizens who kill a relative or acquaintance in a moment of anger only because a gun was available. In fact, every study of homicide shows the overwhelming majority of murderers are career criminals, people with lifelong histories of violence. The typical murderer has a prior criminal history averaging at least six years, with four major felony arrests.

The text is clear and we should not have any trouble in following it. However, upon closer examination, we can ask a few important questions:

1. What does "*banning guns*" mean? Do we mean "all guns" or rather some types of guns?
2. What does "*most murderers*" mean? More than 50 percent? Do we distinguish between murderers, or rather, do we put them all into one group?
3. What does "*overwhelming majority*" mean? 90 percent? more?
4. What does "*typical murderer*" mean? How do we classify murderers so that we emerge with a "typical" pattern?

We should see a connection now between **Rule #1** and critical thinking. Many of the "right questions" that are essential during the critical thinking process are the same questions we ask while looking at the puzzles in this chapter:

- Which words or phrases are ambiguous?
- What are the assumptions?
- What significant information is omitted?

Keeping this firmly in mind, we are now ready to move to the next chapter, where we show that our intuition is often a bad advisor and that we should not jump to conclusions too quickly.