

Teaching Puzzle-based Learning: Development of Basic Concepts

NICKOLAS FALKNER, RAJA SOORIAMURTHI and
ZBIGNIEW MICHALEWICZ

Abstract. While computer science and engineering students are trained to recognise familiar problems with known solutions, they may not be sufficiently prepared to address novel real-world problems. A successful computer science graduate does far more than just program and we must train our students to reach the required levels of analytical and computational thinking, rather than hoping that it will just ‘develop’. As a step in this direction, we have created and experimented with a new first-year level course, Puzzle-based Learning (PBL), that is aimed at getting students to think about how to frame and solve unstructured problems. The pedagogical goal is increase students’ mathematical awareness and general problem solving skills by employing puzzles, which are educational, engaging, and thought provoking. We share our experiences in teaching such a course – apart from a brief discussion on our pedagogical objectives, we concentrate on discussing the presented material which covers (in two lectures) just one selected topic (pattern recognition). In this paper we present the ideas behind foundations for PBL and the material of the first of two lectures on pattern recognition, in which we address core concepts and provide students with sufficient exemplars to illustrate the main points.

Key words and phrases: puzzle-based, computational thinking, problem-based.

ZDM Subject Classification: A20, B50, B70, D40, D50.

Introduction

Students often have difficulties in independent thinking or problem-solving skills regardless of the nature of a problem. At the same time, educators are interested in teaching “thinking skills” rather than “teaching information and content.” The latter approach has dominated in the past. As Fisher (2001) wrote:

“... 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.” Further, many analysts lament the decreasing mathematical skills of students. A recent *Mathematics Working Party Final Report*, issued by the University of Adelaide (June, 2008) includes statements such as “There is an urgent need to raise the profile and importance of mathematics among young people. . . .” and “The declining participation in mathematics and related subjects is not limited to Australia. . . .”

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 are the consequence of decreasing levels of mathematical sophistication in modern societies (see also the book of John Allen Paulos, *Innumeracy: Mathematical Illiteracy and Its Consequences*). Thus it is necessary, to connect thinking and problem-solving skills with mathematical awareness as the current approaches are not satisfactory.

There is no question that these skills are valued in industry and academic circles. A number of large and important employers of computer science graduates value problem and puzzle solving skills as an indicator of the potential employee, including Google, Microsoft and Yahoo (Poundstone, 2000). Industry is full of examples of under-specified problems that, with good analytical and thinking skills, can become correct programs and produce valid results. If we do not train our graduates to think, to be able to apply all of their thinking skills to a problem, we have, ultimately, not properly equipped them for their future.

We believe that a new approach is needed to convey these skills. To address this gap in the educational curriculum, we have created a new course that is aimed at getting students to *think* about how to frame and solve *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 (*e.g.*

How to Solve It: A New Aspect of Mathematical Method) and Martin Gardner during the last 60 years. In one of our favourite 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.”

In this (and the follow-up) paper we share our experiences in teaching such a course – apart from a brief discussion on our pedagogical objectives, we concentrate on discussing the presented material which covers (in two lectures) just one selected topic (pattern recognition). Despite the fact that such discussion is based on a single topic presented during the whole course, still it should provide detailed information on how the puzzle-based learning courses are taught. In this paper we present the ideas behind foundations for PBL and the material of the first of two lectures on pattern recognition, in which we address core concepts and provide students with sufficient exemplars to illustrate the main points. The follow-up paper presents the material of the second of two lectures, in which additional exercises are discussed to reinforce the lesson. The follow-up paper discusses also the outcomes of PBL courses, which include expected improvement in the overall results achieved by students who have undertaken PBL courses, compared to those students who have not.

The Puzzle-based Learning Approach

Today’s marketplace needs graduates capable of solving real problems of innovation and a changing environment – clearly we need more skilled graduates. What is missing in most curricula is coursework focused on the development of problem-solving skills. Most students never learn how to think about solving problems in *general* – throughout their education, they are constrained to concentrate on textbook questions at the back of each chapter, solved using material discussed earlier in the chapter. This constrained form of “problem solving,” is not sufficient preparation for addressing real-world problems. On entering the real world, students find that problems do not come with instructions or guidebooks. One of our favourite examples to illustrate this point is a puzzle on breaking a chocolate bar:

A rectangular chocolate bar consists of $m \times n$ small rectangles and you wish to break it into its constituent parts. At each step, you can only pick up one piece and break it along any of its vertical or horizontal lines. How should you break the chocolate bar using the minimum number of steps (breaks)?

If you do not know the answer, which textbook would you search to discover the solution? Textbooks on optimisation? Simulation? Strategies? Games? Other textbooks? Or it might be that someone wrote a book on chocolates where in chapter 7 there is a full discussion on efficient breaking strategies of a chocolate bar? Very unlikely. The same applies to solving many real world problems: which textbook should you search to find a solution, if that is the solution strategy that you've learned?

Many teachers have used puzzles for teaching purposes and the puzzle-based learning approach has a much longer tradition than just 60 years (Danesi, 2002). The first mathematical puzzles were found in Sumerian texts circa 2,500 BC. However, some of 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 artificial intelligence textbooks to educate computer science students.¹

For additional information on puzzle-based approach the reader is referred to the textbook for this course: *Puzzle-based Learning: An introduction to critical thinking, mathematics, and problem solving* (Michalewicz & Michalewicz, 2008) and/or to earlier articles on this topic (Falkner, Sooriamurthi, Michalewicz, 2010) and (Falkner, Sooriamurthi, Michalewicz, 2009). However, in essence, the puzzle-based learning courses just concentrate on *educational puzzles* that support problem-solving skills and creative thinking. These educational puzzles satisfy most of the following criteria:

1. *Independence*: The puzzles are not specifically tied to a single problem-solving domain.

¹The puzzle is the “river crossing problem”: *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. All the same, the man carries wolf, goat, and cabbage across the river. How has he done it?*

2. *Generality*: Educational puzzles should explain some universal mathematical problem-solving principles. This is of key importance. Most people agree that problem solving, like any other skill, can only be learned by deliberate practice, i.e., 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.
3. *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, including the universal mathematical problem-solving principles, is also remembered.
4. *Eureka factor*: Educational puzzles should initially frustrate the problem-solver, but with the promise of resolution. A puzzle should be interesting because its result is counter-intuitive: problem-solvers often use intuition to start their quest for the solution and this approach can lead them astray. Eventually a *Eureka!* moment is reached (Martin Gardner's *Aha!*), when the correct path to solving the puzzle is recognized.
5. *Entertainment factor*: Educational puzzles should be entertaining and engaging. Entertainment is often a side-effect of simplicity, frustration, the Eureka factor, and an interesting setting.

Educational puzzles can play a major role in attracting students to programs in computer science and mathematics and can be used in talks to high school students and during open-day events. Puzzles can also be a factor in retaining and motivating students. Further, there is a strong connection between the ability to solve puzzles and the ability to solve industry/business problems (Poundstone, 2000).

Puzzle-based Learning vs. Problem-based Learning vs. Project-based Learning

As has already been discussed in (Falkner, Sooriamurthi, Michalewicz, 2010), the ultimate goal of puzzle-based learning is to lay a foundation for students to be effective problem solvers in the real world. At the highest level, problem solving in the real world calls into play three categories of skills: dealing with the vagaries of uncertain and changing conditions; harnessing domain specific knowledge and methods; critical thinking and applying general problem solving strategies. These three skill categories are captured in the three forms of learning depicted below:

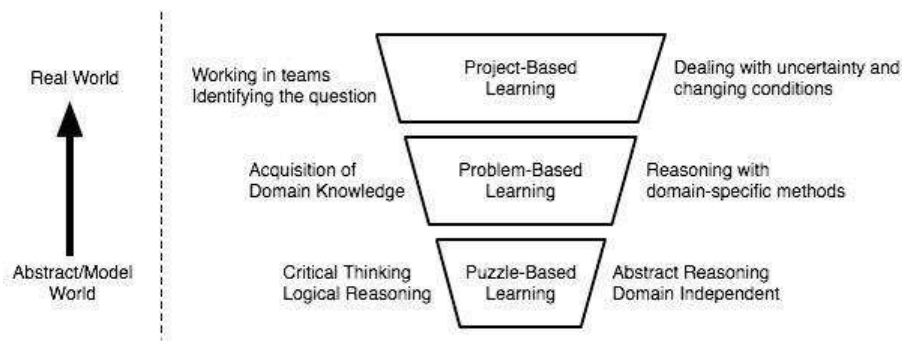


Figure 1. A continuum of learning and skills needed for problem solving in the real world. In this continuum, each layer of skills builds upon the layers below it. The focus of puzzle-based learning is on domain independent, transferable skills. In addition, we aim to foster introspection and reflection on the personal problem solving process. What was I thinking? What is the solution? Why did I not see it?

Both problem-based learning and project-based learning are well established methodologies (Blumenfeld et al. 1991, Bransford et al. 1986). By our description above, problem-based learning requires significant domain knowledge. This is the form of learning one typically sees emphasized in a domain specific undergraduate course such as electromagnetism, data-structures, or circuit-theory etc. Project-based learning on the other hand deals with complex situations where usually there is no one clear unique or correct way of proceeding, for example: How can we increase the adherence of cystic-fibrosis patients to follow their treatment protocol? It can be very hard to determine the best solution. The pedagogical objectives of project-based learning include dealing with ambiguity and complexity, integration of a variety of approaches, user-testing of the value of proposed solutions, and working with a team of people with diverse backgrounds and skills. In both problem-based and project-based learning the problem drives the learning: students need to assess what they already know, what they need to know to address the problem, and how to bridge the knowledge/skill gap.

Puzzle-based learning focuses on domain independent critical thinking and abstract reasoning. This leads to the question: What is the difference between a puzzle and a problem? One way of characterizing the difference is the extent to which domain specific knowledge is needed to solve it. The general flavour of puzzles is that their solution should only require domain neutral general reasoning skills – a biologist, a musician, and an artist should all be able to solve the same

puzzle. The different styles of reasoning required for problem-based learning and puzzle-based learning could be compared to the difference between an in the field investigator and an armchair detective – one only requires reason.

Puzzle-based learning shares many of the pedagogical goals of the emerging paradigm of Computational Thinking (Wing 2006). Puzzle-based learning resonates with the Computational Thinking emphasis on abstraction and analytical thinking. With reference to (Figure 1), Computational Thinking straddles the whole problem skill spectrum but places more emphasis on Problem-based and Project-based learning. With its emphasis on domain independent, rigorous and transferable reasoning we believe that puzzle-based learning lays a basis for CT in the curriculum.

The Puzzle-based Learning Course

There are a few different versions of the puzzle-based learning course available. The course can be offered as a full-semester (three units) elective course or freshman seminar (typically 3 contact hours per week), one unit freshman seminar or one unit core module as part of other courses. Despite a variety of possible offerings of puzzle-based, the structure of the course is very much the same. The topics listed below correspond to weeks of a 12-week semester:

1. **Introduction:** What it is all about?
2. **The problem:** What are you after?
3. **Intuition:** How good is it?
4. **Modeling:** Let's think about the problem
5. **Some mathematical principles:** Do you see it?
6. **Constraints:** How old are my children?
7. **Optimization:** What is the best arrangement?
8. **Probability:** Coins, dice, boxes, and bears
9. **Statistically speaking:** What does it mean?
10. **Let's simulate:** Can we generate the answer?
11. **Pattern Recognition:** What is next?
12. **Strategy:** Shall we play?

Earlier work discussed the general concepts behind puzzle-based learning (Falkner, Sooriamurthi, Michalewicz, 2010) and first experiences in offering such courses


(Falkner, Sooriamurthi, Michalewicz, 2009). In this paper we wish to provide a more practical guide to presentation and concentrate on presenting full material for covering a single topic from the above list of topics. We assume that two classes (50 minutes of instruction time per class – for other lecture lengths, the material can easily be adopted by adding or dropping some slides) is available. Such discussion might be useful for understanding the important points about puzzle-based learning courses, i.e., that the course is not about presenting and discussing a variety of puzzles but rather about presenting, discussing, and understanding problem-solving principles and some mathematical principles in the context of puzzles that serve as entertaining illustration of the presented concepts. Also, the process of understanding problem-solving principles leads students through a variety of topics, exposing them to many important concepts at early stages of their college education.

Topic #11: Pattern Recognition

For this illustrative “lecture”, we have selected topic #11: *pattern recognition*. During the class we consider pattern recognition in the specific sense where we look at sequences of symbols, numbers, actions, and events, and try to discover the pattern in such sequences. The point is that once we identify the pattern, it might be easier to predict the next (or missing) symbol, number, action, or event (in the same way that fraud detection systems try to discover patterns in historical data and then use these patterns to predict which new transactions might be fraudulent). Further, in discussing some issues related to pattern recognition we also talk about many additional topics, including mathematical induction, calculus of finite differences, Diophantine equations, Fibonacci numbers, golden ratio, some problem-solving strategies. This is one of the highlights of the puzzle-based learning courses – starting from one “theme” topic, we diverge into many issues relevant for students in computer science and mathematics.


The benefit of using pattern recognition, as part of a puzzle-based learning course, is that there are so many practical examples to draw from and we raised the issue of fraud detection earlier. Credit card fraud detection revolves around establishing a normal purchasing and usage pattern for a given customer’s card and then detecting significant and unexplained deviations from this pattern. We are required to both identify the original pattern and provide a difference template that will register an anomaly when it corresponds to criminal activity. Still within the financial sphere, consider posing the following question to students:

“The following graph provides daily sales data for a given product from a large food and beverage chain. Where’s the pattern? Do you have any idea what this product is?”



Initial approaches to a solution will probably ask for more information, some of which will be important and some of which will be trivial. For example, the nature of the quantities on the y axis is not that significant – what is important is that, on some days, much more of this product is sold than others. Students should now start to seek other patterns: is there a weekly, monthly or quarterly pattern? They should eliminate ‘obvious’ products with purchasing cycles, such as food or beverage items associated with holidays or religious festivals as the quantities never drop away far enough to reflect the nature of these products.

This puzzle does have a pattern solution but there is not yet enough information to give a good solution to this problem. The trend does seem to be dropping from the February period and this may give some insight as to what is happening here. An astute student may now ask a question such as ‘Which hemisphere or country is this data from?’ If the answer ‘the southern hemisphere’ is given, where summer is from December to February, it appears that whatever is being sold sells more frequently at hotter times of the year. If we then add the missing temperature data (given in blue below) that accompanies the original sales data, we now see a strong pattern:



The correlation² is very strong and not very surprising, when you realise that the product being sold is ice cream. On hot days, people buy far more ice cream. What is interesting, from a sales perspective, is that ice cream is still sold even on cold days!

This problem/puzzle is important because we can illustrate the power of the pattern as a general principle for finding solutions and also convey to the students that sometimes the pattern is not obvious because they do not yet have enough information to see it. This can also be used as a motivator for the students to undertake their own research to solve the problem.

Another excellent question to teach the importance of pattern recognition using puzzle-based learning is to ask the students “You believe that an accountant has been changing company data to hide fraud. How would you detect this if you cannot compare the original and the altered records?” We will discuss the solution to this towards the end of the paper.

Having justified our choice of the pattern recognition topic, the following section presents such lecture material in a narrative way – we do believe that this style of the presentation is better than a reproduction of lecture slides with some comments.

²Earlier in the course when discussing the use of statistics in problem solving we highlight the critical difference between correlation and causation. This example provides a context to revisit these concepts.


A Sample Lecture - Part 1 of 2

We start the class by emphasizing that our ability to recognize patterns is of utmost importance. If we can identify a pattern, then we can build a model to find a solution (e.g. to find the next occurrence of a symbol, number, action, or event). Marilyn Burns, in her book *I Hate Mathematics*, wrote: “*The password of mathematics is pattern.*” Indeed, in many branches of mathematics we search for patterns which allow some generalizations. But we search for patterns everywhere; we even recognize patterns in words:³

Aoccdrnig to a rscheearch at Cmabrigde Uinervtisy, it deosn't mtttaer in waht oredr the ltteers in a wrod are, the olny iprmoetnt tihng is taht the frist and lsat ltteer be at the rghit pclae. The rset can be a toatl mses and you can sitll raed it wouthit porbelm. Tihs is bcuseae the huamn mnid deos not raed ervey lteter by istlef, but the wrod as a wlohe.

(The class usually enjoys the slide with the above statement). Many real-world problems require the prediction of a variable's value (for tomorrow, next month, etc.). A classic example would be to predict sales of a product for the next quarter or year. Plenty of historical data are usually available and the prediction model⁴ may include many economic indicators as variables (e.g. employment, financial, survey, production, and sales indicators). In some problems, however, it is necessary to search for patterns where the data stream is somewhat “irregular” (as opposed to a sequence of numbers recorded every quarter). In such cases it is much harder to discover the pattern, as the following puzzle illustrates (this is the first puzzle-slide displayed during this lecture; the students have a minute or so to think about the solution before the lecture continues):

PUZZLE 1. *The following sequence of seven symbols (commonly known as the M-heart-8 sequence) is “meaningful” in the sense that it is not random:*



What is the next symbol in the sequence?

³This text circulated on the Internet in September 2003.

See: <http://www.mrcctu.cam.ac.uk/~mattd/Cmabrigde/> for translations in other languages.

⁴A *prediction model* is a model that can predict the value of a variable on the basis of the values of other variables.

There are many ways to analyse this sequence (i.e. many ways to search for a pattern)⁵. One possibility is to analyse the features of each symbol. For example, the first symbol consists of four line segments, whereas the second symbol consists of one line segment and two curves. Following the occurrences of line segments and curves, we notice that the symbols consisting of only line segments appear on the first, fourth, and seventh positions. Does this mean that the next symbol consisting of only line segments will appear on the tenth position? If so, can we find a pattern in the number, length, and position of these line segments? This is much harder, because:


- The first symbol consists of four line segments: two long and two short; the two long segments are vertical and the two short segments run at 45 degree angles.
- The fourth symbol consists of five line segments: three long and two medium; of the three long segments, two are vertical and one is horizontal, whereas the two medium segments run at 45-degree angles.
- The seventh symbol consists of three long line segments: one horizontal and two at angles that are greater than 45 degrees.

Even if we are convinced that the tenth symbol consists of only line segments, it would be impossible for us to determine the number, length, and orientation of these segments. Furthermore, it would be even harder for us to analyse the curves of the second, third, fifth, and sixth symbols!

Note that mathematical notation of a sequence is:

$$s[1], s[2], s[3], \dots$$

where $s[i]$ indicates the i -th symbol in the sequence. Thus, in any sequence there is a clear correspondence between the symbols and natural numbers (i.e. the first symbol, second symbol, third symbol, etc.). In the above case, if we number all the symbols:



we will immediately notice that the sequence represents the initial sequence of natural numbers 1, 2, 3, etc. such that each number is displayed alongside its

⁵Interestingly, it has been empirically observed that children are better able to solve this puzzle than adults.

mirror image. With this observation, we should have no difficulty drawing the next symbol! This puzzle is now well known in popular culture due to its placement in an episode of the Fox Network's cartoon series "The Simpsons". In that episode, much was made of a character's inability to solve the puzzle, when a number of other people around her could apparently solve it easily. This puzzle also appeared in the book/film "The Oxford Murders."

In this puzzle, it was not any particular feature (or group of features) of the sequence that led us to the solution⁶, but rather, some external knowledge and our ability to connect this knowledge with the sequence. This was the reason why this puzzle is relatively difficult, as it is unclear what type of knowledge is appropriate or desirable. The second puzzle illustrates the same point, namely, that external knowledge comes in different shapes and sizes (again, the students have a minute or so to think about that before the lecture continues):

PUZZLE 2. What is the missing letter (marked by the "?") in the sequence:

A ? D F G H J K L

And again, we can analyse this sequence by characterizing each symbol by its features. For example, the letters **A**, **F**, **H**, and **K** each consist of 3 line segments. The letter **L** consists of only two line segments, whereas letters **D**, **G**, and **J** include some curves. Is there any pattern this? Or is it important to be more specific and distinguish between longer and shorter segments? For example, the letters **A**, **F**, **H**, and **K** consist of 3 line segments: two long and one short. Is that useful?

Or it might be that we need a different approach based on numbers (which was useful in solving the previous puzzle). There is an obvious correspondence between letters and numbers (as **A** is the first letter, **B** is the second, etc). So, we can translate the sequence in question into the following sequence of numbers:

1 ? 4 6 7 8 10 11 12

This is a growing sequence and the growth is a pattern. If we believe that the pattern is genuine (i.e. it did not arise by chance), then we can conclude that the second number must be **2** or **3**, thus the missing letter is **B** or **C**. But which of these two?

Actually, none of them... The "obvious" answer is: **S**, as the sequence

A S D F G H J K L

⁶In this case, the "solution" is the identification of the pattern.

is the middle row of letters on a computer keyboard!

These two puzzles illustrate the fact that many pattern recognition activities are based on *a priori* knowledge. In general, pattern recognition activities can be based either on information extracted from a sequence (like the fraud detection example, where a system finds patterns in historical data and uses these patterns to predict which new transactions might be fraudulent) or on some *a priori* knowledge. The first two puzzles were difficult as the information extracted from the presented sequences was not helpful, and it was not clear what type of knowledge should be applied. Furthermore, pattern recognition is often related to prediction (as we saw in puzzle 1), where the task was to predict the next symbol in the sequence.

Puzzle 2 is a very interesting puzzle to apply in the contemporary classroom as we have observed students who look at the puzzle and then go back to their (open) laptops. There is often a moment of mild surprise and then realisation as they discover that the answer is right in front of them. This introduces another point, the importance of context for information retrieval.

One of the authors introduced a warm-up exercise for puzzle-based learning activities that requires the participants to think of ten three-letter words for body parts – no slang or ‘rude’ words. Most individuals will be able to manage from five to seven. If we provide some framing context and give the hint that “five are above the neck and five are below the neck”, we may obtain two or three more. We generally only achieve all ten when we combine the solutions across a large group, more than twenty people. Why is this puzzle so challenging? Everybody, when presented with the answers, immediately sees the words that they have missed. If we had pointed to a part of the body and asked for the name, we would have had an immediate answer, yet, asking for the information in a relatively context free environment limits the ability to solve the puzzle.

How do we define context? Are we discussing the cognitive context that frames the knowledge or the physical context that frames the puzzle-solving experience? Students in our courses learn relatively quickly that a seemingly simple puzzle may not have the obvious answer, when asked in one of our lecture theatres – this is the physical context for puzzle solving. To master the broadest application of the knowledge that students have, they must be able to work in an undefined or limited cognitive context and to be able to bring their skills to bear regardless of where they are solving the problem, in an undefined physical

context. While the role of context on problem solving has been documented⁷ one of our core ambitions is to produce thinking students who can solve the puzzles and problems that confront them without having to be given explicit context!

As a final illustration of the constrictive nature of context, consider the recitation of the alphabet. Every student can recite the alphabet – how many can spontaneously reverse it? Have they learned the position of each letter or have they placed each letter into a highly specific contextual framework (possibly with a musical theme)? How many students can tell you what the 16th letter of the alphabet is? While it is possible to derive this information, students do not necessarily have this knowledge at their finger tips – which leads to mistakes and hesitation when they are asked for that information.

The alphabetical puzzle 2 is also a very convenient way to introduce methods for solving puzzles. Polya's heuristics were proposed as a set of approaches to consider in order to find a way to break into the puzzle and start the solution process. We have deliberately moved away from these, although we certainly do not question their worth, as we were concerned that students would spend more time memorising heuristics than learning to solve puzzles. Instead, we use puzzles to illustrate key points and to provide reinforcement for using a certain approach.

This is where puzzle 2 is very useful. Many students guess B or C as mentioned previously and this appears sound. The sequence appears to be strictly increasing and this, as well as the small number of available letters between A and D, reinforces the idea that the answer must be B or C. While looking for ascending or descending patterns, with or without repetition, can be an excellent way to solve a puzzle like this, it doesn't always work. By using this puzzle, we introduce several good approaches for puzzle solving and remind students to use all of the knowledge that they have, even if it is being used out of context.

When we search for patterns in sequences and the search is based on information extracted from a sequence, it is often useful to start with a short sequence and gradually make it longer and longer to discover the pattern. The next puzzle illustrates this process in terms of a number of elements in a set (rather than a sequence):

PUZZLE 3. There are 2^n tennis players, but two of them are twins. The typical tournament rules apply: There are 2^{n-1} games in the first round. The winners advance to the second round, which consists of 2^{n-2} games, and so forth.

⁷See DeFranco and Curcio (1997) for a study on the role of an authentic setting to aid children's problem solving.

Assuming that each player has a 50 – 50 chance of winning any game, what is the probability that the twins will play each other at some stage in the tournament?

To help frame the problem, we briefly introduce the tennis ‘draw’, where players are chosen for each round. We start by assuming that we have enough players for at least one match – that is, 2, i.e. $n = 1$. In the first round, we randomly allocate pairs chosen from the available set. These pairs then play each other and a single winner emerges from each match. Until we have a winner, this number must be even, and we then allocate pairs from this smaller set until everyone has an opponent. Once we have only 1 player remaining, we have a winner.

The formulation of this problem is quite clear – the task is to find the probability p of the meeting of two twins in the tournament in the tournament, which is a function of n (the number of players is 2^n). But how can we calculate this probability? Let us start with smaller sets of players (i.e. small values of n) to see if a pattern would emerge:

- If $n = 1$ (i.e. there are $2^1 = 2$ players in the tournament), then $p = 1$ as they will meet for sure.
- If $n = 2$ (i.e. there are $2^2 = 4$ players in the tournament), then $p = 1/2$. This is because there is a $1/3$ chance that the twins will be paired together in the first round, and a $2/3$ chance that they will play other opponents in the first round. In the latter case, there is a $1/4$ chance that they will meet in the second round (as both of them must win their first round games), and there is a $1/2$ chance for each of them to do that. So:

$$p = 1/2, \text{ because } 1/3 \times 1 + 2/3 \times 1/4 = 1/2$$

- If $n = 3$ (i.e., there are $2^3 = 8$ players in the tournament), then we have three sub-cases to analyse:
 - i. The twins play each other in the first stage. Given that there are a total of 8 players, the probability that the twins will be paired in the first round is $1/7$
 - ii. The twins play each other in the second stage. Here we have two events we need to account for (a) the twins should not have played each other in the first stage and (b) the twins meet in the second stage. For both of these to happen the twins need to be in the first half in the first round, and not have played each other, $2/7$, and each twin needs to have won their game ($1/2 \times 1/2$). So the total probability is $2/7 \times 1/2 \times 1/2 = 1/14$

- iii. The twins play each other in the third stage. Here we three events to account for: (a) the twins should not have played each other in the first stage (b) they should not have played each other in the second stage and (c) they play each other in the third stage. For (a) and (b) to occur the twins need to be different halves during the first stage ($4/7$). For (c) to happen, each twin has to win stage 1 and stage 2 which is $(1/2 \times 1/2) \times (1/2 \times 1/2) = 1/16$. So probability for (iii) is $4/7 \times 1/16 = 1/28$.

Adding the probabilities for these three we get $1/7 + 1/14 + 1/28 = 1/4$ (During class room discussions we draw a decision tree which helps to figuratively convey the pairings and outcomes discuss above when $n = 1, 2$, or 3 .)

So:

$$\begin{aligned} p(1) &= 1, \\ p(2) &= 1/2, \quad \text{and} \\ p(3) &= 1/4. \end{aligned}$$


Thus, a reasonable assumption is that $p(n) = (1/2)^{n-1}$. Indeed, this is accurate and can easily be proved by the induction principle (it depend on the type of students in the class whether we present a simple proof or not).

This illustrates one of the great benefits of puzzle-based learning: providing an easily understood rationale for using a technique. While it would be easier for all lecturers if students learned new knowledge out of pure love or merely because we asked them to, we cannot escape from the importance of motivation.

The Tennis Player puzzle neatly provides us with an example of the importance and power of induction, a theme that is repeated throughout pattern recognition, as we will see below.


The process of stating the right hypothesis is not always that simple. The following puzzle illustrates another interesting point: Sometimes it is difficult to estimate how many smaller sequences should be analysed to reach the correct hypothesis.

PUZZLE 4. *n points are placed on a circle and every point is connected by a line to every other point. Into how many pieces is the circle divided? The following figures illustrate the five initial cases (for $n = 1, 2, 3, 4, 5$):*



The following table would lead us to believe that the number of pieces grows as 2^{n-1} , so adding the 6th point should result in 32 pieces. In actuality, we only get 31 pieces when adding the 6th point:

n	Number of pieces
1	1
2	2
3	4
4	8
5	16



and the right formula for the number of pieces is:

$$(n^4 - 6n^3 + 23n^2 - 18n + 24)/24$$

So we have to be careful not to jump to conclusions too quickly. This also introduces a key point that we have used in lectures – the illustration of a principle through the use of a specific puzzle. Here, we use pattern recognition both as a tool to inspire thought, but also to provide a salient warning – ensure that you have seen the entire pattern before you decide what the pattern is. This is also an opportunity to identify the power of induction, as an inductive proof can establish that the pattern will apply to all future n , rather than the possibility that we have not seen a sufficiently large n to break the pattern!

If there is additional time during the lecture, this is a good moment to introduce the calculus of finite differences, with a few appropriate exercises (e.g. lines cutting a plane) and a brief discussion on Newton's formula.

If the course on puzzle-based learning includes additional time (e.g. tutorial), it is recommended that the students watch (with the instructor) a video of famous Polya's lecture "Let us teach guessing" wherein Polya beautifully illustrates several problem solving heuristics (that are embraced by puzzle-based learning) in the process of deriving a solution to the 5-plane problem. As in the puzzle 4, the maximum numbers of segments in the 3-dimensional space generated by zero, one, two, or three planes, are 1, 2, 4, and 8, respectively, again suggesting a (wrong) pattern.

As human beings, we are great at seeing patterns – so good, in fact, that we can invent them when they do not exist! We draw students' attention to a parallel between the last puzzle to some cases where misleading pattern is created on purpose (e.g. to cover a crime, for counter-intelligence, for competitive situations in the marketplace, or just for fun). One well-known example of this takes place in the short story "*The ABC Murders*" by Agatha Christie. In the story, a murder is committed in Andover and the victim's name was Alice Ascher. Four days later, another murder is committed in Bexhill-on-Sea and the victim was Betty Barnard. Four days after that, Carmichael Clarke was murdered in Churston. The point of the story was that the murderer's goal was just to kill Carmichael Clarke; the purpose of the previous two victims was to create a misleading pattern (useful for avoiding an investigation on the motive of the main crime).

Another enjoyable example of misleading circumstances (again, students are really interested in such stories, which also break a flow of the class material) is credited to José Raúl Capablanca, a Cuban-born world-champion chess player (1921–1927). Apart from being referred to by many chess historians as the *Mozart of chess*, he also displayed an unusual sense of humour. One day he was on a train that was stopped in the middle of nowhere, and he was told it would take many hours to clear the tracks and continue the journey. While Capablanca was waiting, he was approached by a rail-man (who did not recognize Capablanca) and invited to a game of chess. Capablanca feigned that he did know how to play chess. The rail-man was not discouraged and offered a quick lesson. After a brief overview of the game, the rail-main said: "*Let's play. Since you are new to the game of chess and I am an experienced player, it would be only be fair if I play without my queen.*" They played ten games in which Capablanca enjoyed an advantage over his opponent, who played without the most powerful piece – his queen. Capablanca deliberately lost all these games. At the end of the tenth game, Capablanca said: "*After these ten games, I think I know what's most important in chess. Let's play another ten games, but this time I will remove*

my queen and you play with a full set.” The rail-man was surprised (to say the least!), but obliged. They played another ten games and this time Capablanca won all ten. He then said to the dazed rail-man: *“I knew it from the start: It is much easier to win chess if I play without a queen!”*

It is not surprising that many puzzles are constructed in such a way as to mislead us into the “wrong” pattern – this is, after all, the nature of the puzzle!

Conclusions

The graduates of tomorrow need the ability to adapt their problem solving skills to problems that are not only beyond what they have already experienced, but are potentially outside of the context within which they were trained. We see puzzle-based learning as the first building block to develop skills that may be applied, regardless of the context or knowledge domain. Computer Science education goes beyond the syntax and semantics of programming languages to provide the framework for the development of mature analysts, graduates capable of considering the problem fully before coming to a reasoned and well-supported conclusion as to how they will move forward in algorithmic development.

Puzzle-based learning is an experiment in progress. The goal is to foster general domain independent reasoning and critical thinking skills that can lay a foundation for problem-solving in future course work. As fun as puzzles inherently are, they are just a means to this pedagogical end. Our preliminary experience in different contexts has been encouraging and well received as we continue to explore this approach. As instructors for this new course it has been a learning experience for us also.

In this paper, we have discussed in detail one foundational critical thinking skill we would like to foster: *recognizing and reasoning with patterns in data*. This component of the course is often spread over two class periods. The first of these two class periods has been discussed in detail in this paper with several illustrative examples. This first lecture presents the student with an introductory set of puzzles in order to engage their interest, start their creative thinking processes and identify a new area in which they can develop a set of appropriate puzzle-solving heuristics. In discussing some issues related to pattern recognition we also talk about many additional topics, including mathematical induction, calculus of finite differences, Diophantine equations, Fibonacci numbers, golden ratio, and additional some problem-solving strategies. This is one of the highlights of the puzzle-based learning courses – starting from one “theme” topic we diverge into many issues relevant for students in computer science and mathematics.

This is not, however, the entirety of the lecture material for this course. The second lecture goes into more detail, provides more worked examples and reinforces the lessons learned. Where two separate lecture slots are not available, we would normally condense the material into one slot, reducing this introductory material to an introductory framework and providing more of the detailed work.

The format, and the nature, of the second lecture are designed to encourage the development of transferable skills in which additional exercises are discussed to reinforce the lesson of taking PBL from puzzle-solving to problem-solving and to a mode where the skills can be applied in other domains of interest. The second lecture is described in more detail in our second article on this subject: Teaching Puzzle Based Learning - Developing Transferable Skills. The follow-up paper (Michalewicz, Sooriamurthi, Falkner, 2012) also discusses the outcomes of PBL courses, which include expected improvement in the overall results achieved by students who have undertaken PBL courses, compared to those students who has not.

For more information on the nature of puzzles and the approaches used in puzzle-based learning, readers are directed to the website associated with the text, www.PuzzleBasedLearning.edu.au.

References

- [1] P. C. Blumenfeld, E. Soloway, R. W. Marx, J. S. Krajcik, M. Guzdial and A. Palincsar, Motivating project-based learning: Sustaining the doing, supporting the learning, *Educational Psychologist* **26**, no. 3&4, 1991, 369–398.
- [2] J. D. Bransford, R. S. Sherwood, N. J. Vye and J. Rieser, Teaching thinking and problem solving: Research foundations, *American Psychologist* **41**, 1986, 1078–1089.
- [3] M. Danesi, *The Puzzle Instinct: The meaning of puzzles in human life*, Indiana University Press, 2002.
- [4] T. C. DeFranco and F. Curcio, A Division Problem with a remainder embedded across two contexts: Children’s Solutions in Restrictive vs. Real-World Settings, *Focus on learning problems in mathematics* **19**, no. 2, 1997, 58–72.
- [5] N. Falkner, R. Sooriamurthi and Z. Michalewicz, Puzzle-Based Learning for Engineering and Computer Science, *IEEE Computer* **43**, no. 4 (April 2010), 20–28.
- [6] N. Falkner, R. Sooriamurthi and Z. Michalewicz, *Puzzle-Based Learning: The first experiences (2009)*, Proceedings of the Twentieth Annual Conference of the Australasian Association for Engineering Education (AaeE 2009), Adelaide, Australia, December 6–9.
- [7] N. Falkner, R. Sooriamurthi and Z. Michalewicz, *Teaching Puzzle-based Learning: Development of Transferable Skills*, 2012, to appear in the Journal of Teaching Mathematics and Computer Science.

- [8] A. Fisher, *Critical Thinking: An Introduction*, Cambridge University Press, Cambridge, UK, 2001.
- [9] M. Gardner, *Entertaining Mathematical Puzzles*, Dover Publications, New York, 1961.
- [10] Z. Michalewicz and M. Michalewicz, *Puzzle-Based Learning: An introduction to critical thinking, mathematics, and problem solving*, Hybrid Publishers, 2008.
- [11] Z. Michalewicz, N. Falkner and R. Sooriamurthi, *Puzzle-Based Learning: An Introduction to Critical Thinking and Problem Solving*, Proceedings of the 41st annual meeting of the Decision Sciences Institute, 2010, 3101–3106.
- [12] J. A. Paulos, *Innumeracy: Mathematical Illiteracy and Its Consequences*, Hill and Wang, New York, 2001.
- [13] G. Polya, *How to Solve It: A New Aspect of Mathematical Method*, Princeton University Press, Princeton, 1945.
- [14] W. Poundstone, *How would you move Mount Fuji?: Microsoft's Cult of the Puzzle – How the World's Smartest Companies Select the Most Creative Thinkers*, Little Brown and Company, 2000.
- [15] University of Adelaide: Mathematics Working Party Final Report, internal document, June 2008.
- [16] J. M. Wing, Computational Thinking, *Communications of the ACM* **49**, no. 3 (March 2006), 33–35.

NICKOLAS FALKNER
SCHOOL OF COMPUTER SCIENCE
UNIVERSITY OF ADELAIDE
ADELAIDE, SA 5005
AUSTRALIA
E-mail: jnick@cs.adelaide.edu.au

RAJA SOORIAMURTHI
INFORMATION SYSTEMS PROGRAM
CARNEGIE MELLON UNIVERSITY
PITTSBURGH, PA 15213
USA
E-mail: raja@cmu.edu

ZBIGNIEW MICHALEWICZ
SCHOOL OF COMPUTER SCIENCE
UNIVERSITY OF ADELAIDE
ADELAIDE, SA 5005
AUSTRALIA
AND
INSTITUTE OF COMPUTER SCIENCE
POLISH ACADEMY OF SCIENCES AND THE
POLISH-JAPANESE INSTITUTE OF INFORMATION TECHNOLOGY
WARSAW
POLAND
E-mail: zbigniew.michalewicz@adelaide.edu.au

(Received July, 2011)