Open-ended Problem Solving Via Punctuated Dialogue

Anthony M. Starfield
Karl A. Smith
Civil and Mineral Engineering
University of Minnesota
Minneapolis, Minnesota

Abstract

Our concern with developing students' engineering-design skills prompted us to start building a foundation of problem formulation-and-modeling skills prior to the senior capstone design project. We have attempted to promote the development of these skills in many ways. For example we have used small-group active-learning strategies, computer-assisted learning, and peer tutoring. As class sizes become larger and larger, it becomes increasingly difficult to monitor students' problem-formulation and modeling performance (and to ask appropriate questions and provide guidance).

We are currently experimenting with expanded problem statements in two classes (Applications of Operations Research for Civil and Mineral Engineering Juniors, and Formulation, Modeling and Analysis for Engineering Problems for Engineering and Science Freshmen). First we hand out the problem statement and then, as the students progress through the problem, we provide additional written material. Students are required to ask questions, make assumptions, build models, and propose solutions at regular intervals. Upon review, and sometimes revision of their response, we provide additional information. The process is punctuated by students reporting and by our commenting on and discussing their progress.

Our practice of providing students with a written narrative on a problem is an extension of the slightly open-ended problems that we have been using with students in our engineering classes for the past ten years. Recently we have extended the approach to a book titled How to model it: Problem solving for the computer age.

Problems with Engineering Teaching I

This paper opens, as do our classes, with a question. Ask yourself "Who is learning in the typical college classroom?" Alternatively, ask yourself, "Who is organizing, summarizing, and presenting?"

STOP AND THINK

In the typical college classroom only the instructor is active. The typical interaction between instructor and student is perhaps described best by this comment from one of our professors, "The problem with lecture is that the information passes from the notes of the professor to the notes of the student without passing through the mind of either one." Students who do not understand the material being presented and yet mechanically write down what the instructor is saying illustrate one of the many weaknesses of the lecture. Other difficulties and pitfalls include[1]:

1. Students who are preoccupied with what happened during the previous hour or with what happened on the way to class,

2. Entertaining and clear lectures that misrepresent the complexity of the material being presented,

3. Students who are isolated and alienated and believe that no one cares about them as persons or about their academic progress.

In order to maximize their achievement, especially when studying conceptually complex and content-dense material, students should not be allowed to be passive while they are learning. One way to get students more actively involved in this process is to incorporate cooperative interaction into the college class so that students have to explain what they are learning to each other. They must listen to and understand each other's point of view, give and receive support from classmates, and help each other dig below the superficial level of understanding of the material they are learning. Without creating for the students a learning environment where they may develop and practice the social skills required to work cooperatively with others, how can we, as college faculty, honestly claim that we have prepared our students for a world in which they will often need to coordinate their efforts with others? Skills that our students develop to work cooperatively in the classroom will be applicable not only on the job, but in community, social, and family situations as well.

This approach is consistent with the current state-of-the-art in college teaching. McKeachie, et. al., summarized the research on instruction as follows in the recent NCRIPTAL report Teaching and Learning in the College Classroom[2]:

> The best answer to the question, "What is the most effective method of teaching?" is that it depends on the goal, the student, the content, and the teacher. But the next best answer is, "Students teaching other students." There is a wealth of evidence that peer teaching is extremely effective for a wide range of goals, content, and students of different levels and personalities. (p. 63)

We have described active-learning instruction methodologies in several of our previous articles--

Structuring learning to achieve the goals of engineering education[3], Educational engineering[4], The nature and development of engineering expertise[5].

Problems with Engineering Teaching II

Students have ample experience solving problems that have a unique answer, that is, problems that have an answer (**the answer**) printed at the back of the book. Although these problems have a certain limited usefulness in the improvement of students' mechanical problem-solving skills, they are not appropriate for developing students' modeling abilities. Appropriate problems to assign for the development of students' modeling skills are those that do not have single, unique answers, but whose answers depend on the problem formulations.

Problem formulation has been neglected in the recent drive to promote students problem-solving and higher-level thinking skills. And "problem finding" (as defined by Mackworth[6] in his article "Originality") has been ignored.

Table 1

Problem Solving and Problem Finding
(Mackworth, 1965)

Definition

Problem solving is the selection and use of an existing program from an existing set of programs.

Problem finding is the detection of the need for a new program by comparing existing and expected future programs.

Objective

Problem solving: To choose correctly between existing programs--in order to select the one program that effectively elicits the required actions from a set of possible responses.

Problem finding: To choose correctly between existing and expected future programs--in order to devise new programs and to realize that one or more of these would be more suitable than any of the existing programs in eliciting the required actions.

Method

Problem solving: Experiment more than thought minimizes the mismatch between the desired and apparent actual states.

Problem finding: Thought more than experiment minimizes the mismatch between the desired and apparent actual states.

Outcome

Problem solving: Success is the discovery of one specific acceptable answer to one well-defined problem.

Problem finding: Success is the discovery of many general questions from many ill-defined problems.

One way to assist students' development of these skills is through requiring them to construct models. "Models" here means not only physical models as we often think of them in reference to "modeling" clay or architectural "models", but also mathematical models, implemented either analytically through a functional relationship or numerically through a discrete function solved using a computer. A model in this sense is a representation of a problem that captures some, but not all, of the features of the problem.

Problems in the real world do not magically appear in a form ready to be solved. They are messy and often not clearly identified or, if identified, the label or identification is often inappropriate. The principal problem is often figuring out what the problem is. In short, real problems (in contrast to text-book problems) are not naturally well formulated. Even after identifying the problem, much iteration is usually required to create a satisfactory solution. Students, however, often think that once they have "solved" the problem, that is, generated an answer, they are finished. And for them, the sooner the better (and they don't reconsider their work unless forced to). The attitude engendered is not favorable to tackling and solving the problems they will encounter in the world, be it in their professional career, their family, or their community.

We have addressed modeling and open-ended problem solving in our articles: Constructing knowledge bases[7], Teaching problem solving: An alternative to case studies[8], and Mastering engineering concepts by building an expert system[9].

Our Classroom Approach

For years we have worked to create better ways of engaging students in problem solving, both by assisting in the development of their problem solving skills, and by making learning more relevant, meaningful and interesting. The following example gives an indication of the type of problem we set in an introductory course for engineering, science and mathematics freshmen. The "ping-pong ball problem" is presented in cooperative-learning, lesson-plan format.

As far as we know the ping-pong ball problem originated at MIT. It came to us via Billy Koen[10].

****

Ping-Pong

Karl A. Smith
University of Minnesota

Subject Area: Problem Solving

Grade Level: High School/College

Lesson Summary: An informal work group task to introduce the concept of constructing a model for solving the problem of how many ping-pong balls could fit in the room. Quick, one-minute estimate by each individual followed by five-minute estimate by pairs.

Instructional Objectives: Students will increase their skill at estimating and constructing models, and will learn the value of time as a resource in problem solving.

Materials: No handout materials are required.

Time Required: Approximately 30 minutes.

DECISIONS

Group Size: Two. Pairs can formulate and solve the problem quickly.

Assignment to Group: Self-selection by students if class is very large and time-saving is important, otherwise random assignment.

Roles: Person with the smallest individual estimate is the Recorder. Person with largest individual estimate is the Prober.

THE LESSON

Instructional Task: "The first task is to be completed individually. Look around the room you are sitting in, then take just 60 seconds to answer the question, 'How many ping-pong balls could you fit into the room?'" After one minute call for answers and write down all that are volunteered.

"The second task is to be completed by pairs. Take five minutes to answer the same question and to develop an explanation of how you arrived at your answer."

Call on each pair for their answer. Point out the difference in spread (variance) in the answers between task 1 and task 2. Usually the spread is much less for the five minute answers.

Ask the pairs, "How did you get it?" "Did you guess?" "Did you construct a representation?" "If so, can you describe your representation?"

Inquire whether or not any pairs tried to make a rough calculation of the number of ping-pong balls. Some may, for instance, have estimated how many balls would fit on the wall at one end of the room, and then multiplied by the number that would fit along the length of the room. Or you may have estimated the volume of a ping-pong ball and the volume of the room, and divided the one into the other.

> If any pair made a volumetric calculation, ask for their model of a ping-pong ball? What was their model of the room?

Check to determine if their representation of the ball was a cube rather than a sphere! Also determine if they represented the room as a large, empty box.

If any pairs developed and used a symbolic representation and introduced a notation, follow-up on it. An example is as follows:

let L be the length of the room,
let W be its width,
let H be its height,
and let D be the diameter of a ping-pong ball.

Then the volume of the room is

$$V_{room} = L * W * H,$$

and the volume of a ball (treating it as a cube) is

$$V_{ball} = D^3,$$

so $$\text{number of balls} = (V_{room}) / (V_{ball}) = (L * W * H) / (D^3).$$

An important question to consider is, "How would you arrive at the best answer you can give to the question, 'How many ping-pong balls could fit in this room?'?" Think about this question and discuss it with your partner before giving an answer.

Do you recommend, "Measure the room and ball accurately"? Or do you think that the best answer is "Fill the room up with ping-pong balls and count them!" Can you think of a better approach than that? Is it worth the effort? How good an answer on the number of ping-pong balls are you willing to accept?

Notice that you cannot really answer the question, "How many ping-pong balls could fit in this room?" unless you are told how good an answer is needed.

Positive Interdependence: One answer and explanation from each pair.

Individual Accountability: Randomly choose one member of the pairs to present.

Criteria for Success: Answer from each pair.

Expected Behaviors: Everyone participates in pair discussion. Each member of pair can explain answer and formulation.

MONITORING AND PROCESSING

Monitoring: Circulate among the pairs to check that roles are being followed (only one person recording and one person probing).

Intervening: Remind pairs that both are expected to participate, and to understand and be able to explain their formulation and solution. Model and coach probing by asking questions.

Processing: Remind groups that every member has two functions: to complete the task and to maintain good working relationships. Ask groups to discuss effectiveness by individually listing things that went well and things the need to be worked on. Whip around groups and list the things that went well and the things that need work.

AUTHOR'S NOTE

This problem is included in the forthcoming book How to model it: Problem solving for the computer age, by Starfield, A.M.; Smith, K.A.; and Bleloch, A.L. The book features problem formulation and modeling through a dialogue with the reader, punctuated intermittently with a request for the reader to perform a task.

****

Our Textbook Approach

Getting students involved in constructing models in the classroom (such as for the ping-pong problem) is an effective means of assisting in the development of their skills for engineering. However, as class sizes increase over 100 it is a challenge to keep all of the students involved. Furthermore, it is difficult to involve students in modeling and problem solving outside of the classroom.

Our challenge was to adapt this classroom approach to helping students develop modeling skills to a book. We hope we have accomplished this through a procedure called "punctuated dialogue." By using this procedure of punctuated dialogue we strive to develop an extended "conversation" with our reader. Also, we encourage each reader to find at least one other person to work through the book with because working together, besides being more fun, involves explaining, describing rationale, and elaborating, all of which have demonstrated effectiveness for fostering the development of skills.

Chapter Two from How to model it is reprinted here to illustrate our approach in the book[11]. This chapter and the book features problem formulation and modeling through a dialogue with the reader, punctuated intermittently with a request for the reader to perform a task.

****

Chapter 2: Time for ping pong?

SOLVING A PROBLEM IN SIXTY SECONDS

Look around the room you are sitting in. Then take just sixty seconds to answer the following question:

'How many ping pong balls could you fit into the room?'

STOP AND THINK

Time's up! What is your answer?

> How did you get it? Did you guess? Did you build a model? If so, can you describe your model?

STOP AND THINK

Sixty seconds is not much time, is it? Perhaps it was just enough time to shrug your shoulders, look around the room, and write down 'Lots!' or 'Thousands!' These would not have been good answers if we had given you more time, but there is nothing wrong with them here.

> What have you accomplished if your answer was 'Lots' or 'Thousands'? What type of model did you use?

You probably did not think you were using a model at all, but you were! You modeled in terms of categories such as [few, some, lots] or [tens, hundreds, thousands]. By drawing a mental picture of a ping pong ball, and looking around the room, you then estimated that the answer belonged in one category rather than another. This is not a useless exercise; it makes a difference whether the answer is 'some' or 'lots'.

Alternatively, you might have tried to make a rough calculation of the number of ping pong balls. You might, for instance, have estimated how many balls would fit on the wall at one end of the room, and then multiplied by the number that would fit along the length of the room. Or you might have estimated the volume of a ping pong ball and the volume of the room, and divided the one into the other.

> If you made a volumetric calculation, what was your model of a ping pong ball? What was your model of the room?

We would bet that your model of the ball was a cube rather than a sphere! And you probably modeled the room as a large, empty box. Is Figure 2.1 or Figure 2.2 a fair representation of your model?

> What other simplifications or assumptions did you make?

Sixty seconds did not give you much time to think about the assumptions you were making, but you were probably fleetingly aware of some of them. Did you wonder, for instance, about whether you could ignore the furniture in the room? Did you assume that you were not allowed to squash or deform the ping pong balls?

SOLVING THE SAME PROBLEM IN FIVE MINUTES

Now take FIVE MINUTES to solve the same problem. This time keep a note of how you go about solving it. If possible, find a partner to work with.

STOP AND THINK

> How did you go about solving the problem this time? Did you use the same procedure but refine your measurements? Or did you use the extra time to take a new approach? Did you change your model?
>
> Did you modify your assumptions? Did you, for example, make a correction for the furniture or the shape of the room?
>
> If you worked with a partner, what were the differences between working together and working alone? Did you share out tasks? Did you start with similar ideas, or did you spend time arguing about the proper approach?

Given more resources (time and people) the chances are that you built a more sophisticated model. If your first model was a rough approximation, you might well have switched to a volumetric method such as the model shown in Figure 2.2.

You might even have used a symbolic representation and introduced a notation. For instance, you might have said:
let L be the length of the room,
let W be its width,
let H be its height,
and let D be the diameter of a ping pong ball.

Then the volume of the room is

$$V_{room} = L * W * H,$$

and the volume of a ball (treating it as a cube) is

$$V_{ball} = D^3,$$

so number of balls $= (V_{room}) / (V_{ball})$

$$= (L * W * H) / (D^3). \quad .... (2.1)$$

WHAT HAVE WE LEARNED ABOUT MODELING?

> Take about ten minutes to make a list of the points that you think these two exercises illustrate. This is not an easy task, so if you have run out of ideas, read on.

STOP AND THINK

We will first make a list of points that we think are important, so you can compare it with your list. Then we will expand on the more important points.

1. Both exercises illustrate the point made in chapter 1 that a model is a partial rather than a complete representation.

2. Even a very rough answer is better than no answer at all.

3. A model that is inadequate under one set of circumstances may be the best that you can do under another set of circumstances. It follows that the design of a model depends as much on circumstances and constraints (of money, time, data or personnel) as it does on the problem that is being solved. It also follows that the assumptions one makes depend on the circumstances in which one solves the problem.

4. A symbolic representation (choosing a notation and building a formula or formulae) is 'clean' and powerful. It communicates, simply and clearly, what the modeler believes is important, what information is needed and how that information will be used.

5. Sometimes one uses models implicitly (without being aware that one is doing so); at other times one consciously or explicitly constructs or uses a model. An EXPLICIT model is an indispensable tool for solving problems and for talking about the solution.

> Did you recognize the last point? Would you agree that it is probably the most important point in the list?

We will expand on point 5 after the next section, which relates mainly to points 1 and 3 above.

the real world, the model world and Occam's razor

The room you are sitting in is the 'real world' of this problem. It is possible that the room has an odd shape, curved walls or a vaulted ceiling. There are almost certainly windows and doors. One wall may be painted white, another blue. There may be pictures hanging on the walls, carpets on the floor and lights suspended from the ceiling. There is furniture in the room - chairs, desks, cupboards, etc. And for some obscure reason somebody wants to fill it with ping pong balls.

Your 'model world', on the other hand, is likely to be more like Figure 2.2: a large box filled with small cubes.

> What is the connection between these two worlds? How do we get from one to the other? Why, for instance, does the model world have no windows? Does it make sense to ask whether the walls of the box are painted white and blue?

A model can be likened to a caricature. A caricature picks on certain features (a nose, a smile or a shock of hair) and concentrates on them at the expense of other features. A good caricature is one where these special features have been chosen purposefully and effectively. In the same way a model concentrates on certain features of the real world. Whenever you build a model you have to be selective. You have to identify those aspects of the real world that are relevant and ignore the rest. You have to create a stripped down model world which enables you to focus, single-mindedly, on the problem you are trying to solve.

William of Occam was a fourteenth century Scottish philosopher who propounded a heuristic in Latin 'Entia non sunt multiplicanda'. Translated literally, this means 'Things should not be multiplied', but in the context of philosophy (and modeling) it means that one should eliminate all unnecessary information relating to the problem that is being analyzed. Since Occam was reputed to have a sharp, cutting mind, this heuristic is known as Occam's razor.

It is sometimes useful to think of the model world connected to the real world by a tunnel or passage-way as in Figure 2.3. The passage-way is guarded by a mythical customs officer wielding Occam's razor. It is his job to make sure that nothing inessential is able to pass from the real world to the model world.

One feature distinguishing a good model from a bad model is the way in which Occam's razor has been used. A bad model is the result of either using the razor too little (letting irrelevant details creep into the model) or using it too much (cutting out essential features of the real world). A good model is one that retains a proper balance between what is included and what is excluded.

> How do we reach that balance? How does the customs officer do his job? Certain decisions are easy. For example, nobody would even think of asking whether the colors of the walls should be included in the ping pong model world. But other decisions are not so easy. Should the furniture be in the model world, or should Occam's razor cut it out? And is a cube a good representation of a ping pong ball?

There are no hard and fast rules. The customs officer does not have a list of forbidden imports. But he does have two important guidelines:

1. The purpose of the model. We defined a model in chapter 1 as a purposeful representation. The customs officer should cut out any aspect of the real world that does not contribute significantly to the stated purpose of the model.

2. Constraints of time, money, personnel and information. The more restrictive these constraints, the more ruthless the use of Occam's razor.

> How did we apply these guidelines in the ping pong problem?

The time constraints were so severe that we almost exclusively invoked the second guideline. Figure 2.1 represents the model we thought of in just sixty seconds. If you think about it, the model world corresponding to that figure is incredibly austere: the room is represented by one wall and a length, while the ping pong ball is represented by a square cross-section or a diameter. Even in the five minute exercise, our model world consisted of only a box and cubes, as in Figure 2.2.

Notice that if we had been given more time we would still NOT have been able to use the first guideline effectively for the simple reason that the purpose of the ping pong problem has not been stated clearly enough. We do not know whether the object of the exercise is to obtain a rough estimate (to the nearest ten thousand, say) or whether a significant prize has been offered for the plan that crams in a maximum number of balls. In the first case the windows, doors and light-fittings would be irrelevant, no matter how much time we had for building the model. In the second case, their inclusion or exclusion from the model world would depend only on constraints.

attempting to define heuristics

In chapter 1 we described a heuristic as a rule-of-thumb and we flagged it as a modeling term that we would explain more fully later on. We used the word again in the above section; it is time we tried to define it.

A heuristic is a plausible or reasonable approach that has often (but not necessarily always) proved to be useful; it is not GUARANTEED to be useful or to lead to a solution. This is just a more formal way of saying that it is a rule-of-thumb. Heuristics are difficult to define, but relatively easy to recognize.

> Occam's razor is not the only heuristic introduced in the previous section. What other heuristics can you identify?

The guidelines to the mythical customs officer are heuristics too. They are heuristics for applying Occam's razor.

> Notice how the discussion towards the end of the last section illustrates the point that a heuristic need not always be useful. Using the objectives of a model to filter out those aspects of the real world that should be included in the model is a very important heuristic, but it is almost irrelevant when time constraints predominate (as in the sixty second exercise).

It is precisely because we cannot guarantee the efficacy of heuristics that modeling is an art rather than a science. We will have more to say about this in the next chapter.

why models are important

Have you noticed that we never once said 'Build a model'? We just presented you with a problem and then suggested that either consciously or unconsciously you were using models to solve it. We then went on to TALK about the models you were probably using.

The point we want to make is that thinking CONSCIOUSLY and EXPLICITLY about models is a crucial part of problem-solving in general.

> Why do you think we have stressed the words CONSCIOUSLY and EXPLICITLY?

Think about what your answer to the ping pong ball problem means to the person who asks for it. It does not help very much to give an answer like 'lots' or even '28,517 balls' without also giving an explanation of HOW YOU GOT THERE. The model you have used is as important as your answer in a problem like this.

> Why?

Because in the time available there is no way that you are going to solve the problem precisely. You are bound to take short cuts. It follows that there is no way to evaluate your answer unless one knows more about the assumptions you made and the short cuts you took. If you are not AWARE of the model you are using, you are not going to be able to COMMUNICATE your answer in a meaningful way.

Come to think about it, if you are not aware of the model you are using, you are not going to be able to communicate with YOURSELF. We all need to have an explicit model (or models) to think clearly about any problem we are trying to solve.

This is why models are so important!

> It is possible that you noticed the role of models in communication when you teamed up for the five minute exercise. Did you ask your partner (or your partner ask you) 'What did you do for the sixty second exercise?' Or did you ask 'What was your answer?'

It was far more important to identify your models than to compare answers. If you compared models and then discussed what to do next, you probably did a much better job of the five minute exercise than if you jumped in and both argued from different premises.

IF WE HAD MORE TIME

what is the best answer we could give?

An important question to consider is "What is the best answer you can give to the question, 'How many ping pong balls could fit in this room?'?" Think about this question and discuss it with your partner before giving an answer.

Stop and Think

> Do you recommend to 'measure the room and ball accurately'? Or do you think that the best answer is 'Fill the room up with ping pong balls and count them!' Can you think of a better answer than that? Is it worth the effort? How good an answer are you willing to accept?

Notice that you cannot really answer this question unless the PURPOSE or OBJECTIVE of the exercise is stated more clearly! We are back to the difficulty we noticed when we were discussing guidelines for the use of Occam's razor: nobody has told us why we are filling the room with ping pong balls.

> But it is always a good heuristic to ask about the best answer one could possibly find!

Thinking about the best possible answer is equivalent to drawing up a wish-list of the things you would like to have in your model world. Asking how good an answer you are willing to accept then helps to prune that wish-list. This is where the time and other constraints become important.

It is good to ask the question about the best answer even when, as in this case, you find you cannot really answer it. At least you are alerted to the deficiencies in the statement of the problem. In this case you should be prompted to go back and ask 'Why do you want to know how many ping pong balls will fit in this room?'

upper and lower bounds

In the sixty second exercise you were so pressed for time that your objective was to find an answer, almost any answer, without considering how good an answer it was. The five minute exercise may have given you more time to be critical about your answer. Certainly, if you had more time it would be prudent to evaluate or put error bounds on your answer.

> If you used a volumetric model such as Figure 2.2, did you pause to consider whether your answer would be too small or too large? Assuming that you know the volume of free space in the room, it will under-estimate the number of balls.
>
> Can you think of a model that would just as quickly over-estimate the answer?

Stop and Think

Suppose you model a ping pong ball as a sphere instead of a cube. If its diameter is D then its volume is

$$V_{ball} = (4*pi/3) * (D/2)^3.$$

Now imagine the spheres somehow packed into the room so that there are absolutely no air gaps between them. Equation 2.1 then becomes

$$\text{number of balls} = (V_{room}) / (V_{ball}) = (L*W*H) / [(4*pi/3)*(D/2)^3]. \quad ...(2.2)$$

Equation 2.2 gives an upper bound to the answer.

> How close are the two bounds?

Notice that if you divide equation 2.2 by equation 2.1 you get

(upper bound)/(lower bound) = 6 / pi,

which is a ratio of nearly 2.

This would have been a useful calculation to make, even if you only had five minutes to solve the problem, because it tells you that the type of model you are using produces answers that could easily be out by 50% or more, if only because you have not properly considered how the spheres will fit or pack together.

If you had more time, this calculation becomes even more significant. Should you spend the extra time looking at packing patterns of spheres, or is the volume of furniture or shape of the room more important? You now have an estimate of how important the packing could be; you can compare it with, for example, an estimate of the volume of the furniture as a percentage of the volume of the room.

Notice, by the way, how a symbolic or mathematical representation can produce unexpected benefits. Comparing equations 2.1 and 2.2 enabled us to estimate the ratio of upper to lower bound WITHOUT MEASURING THE DIMENSIONS OF EITHER A BALL OR THE ROOM.

comparing assumptions

Making assumptions is an integral part of deciding what to take from the real to the model world. It follows that assumptions are subject to the two guidelines that help the customs officer to use Occam's razor intelligently. The assumptions one makes depend on both the purpose of the model and the constraints under which it is built.

Asking what we would do if we had more time illustrates how our assumptions depend on constraints. Let us make a list of some of the more important assumptions implicit in Figure 2.2:

- the room is shaped like a box;
- a ping pong ball is assumed to be a cube;
- furniture in the room can be ignored;
- window and door spaces and other nooks and crannies can also be ignored.

Given a little extra time, which of these assumptions should we relax? Should we try to estimate the volume of the furniture? Or improve our model of the shape of the room? Or should we investigate the way in which spheres can be packed together?

If we had a lot of extra time we would probably do all three. If we have only a little time we should invest a part of that extra time RANKING the assumptions. We already (from the section on upper and lower bounds) have some idea of the effect of packing; we should also be able to make rough estimates of the effect of furniture and the shape of the room.

Notice that as we relax the resource constraints, so we can afford to consider more and more obscure assumptions. (What about the space under the furniture?) Each assumption we investigate in turn opens up new assumptions at a finer level of detail. Notice too that each time we address an assumption, we make our model world more detailed, and each detail will require more and more data in the form of measurements (such as the dimensions and shape of each piece of furniture). Eventually we HAVE TO establish the objective of the exercise and ask whether the detail is really necessary.

In other words, we have to determine the appropriate RESOLUTION of the model.

packing spheres

Suppose we decided that the way in which the balls packed into the room was important. How much difference would that make?

In the section on 'upper and lower bounds', we considered two extreme cases of how the ping pong balls might pack into the room. The first was our original model (a ball is a cube) where we ignore the possibility that rows of balls might fit together in a 'tight' pattern. In the second model we assumed that the balls packed so tightly that no air space whatsoever was left between them. The two models respectively give lower and upper bounds to the number of balls that will fit in a box, and not surprisingly, these bounds are far apart.

To get a better idea of how the balls might actually pack, consider Figure 2.4. In Figure 2.4(a) we ignore packing. This is our 'ping pong ball is a cube' model. Notice that the vertical distance between the centers of two balls is just D, the diameter of a ball.

Figure 2.4(b) is a more realistic representation. In this case we have denoted the vertical distance between the centers of two balls by H.

What is the relationship between H and D?

Consider the triangle PQR that joins the centers of three balls. Notice that it is an equilateral triangle; the length of each side is equal to D. It follows that all three angles of the triangle are equal to 60 degrees and hence that

$$H / D = \sin(60)$$

or

$$H = 0.866 D.$$

Figure 2.4 thus suggests that we could pack nearly 14 percent more balls into the room because of the way they fit.

However, notice that the figure is only a two-dimensional representation of a ball. It is often easier to work with two-dimensional rather than three-dimensional representations and it is also often useful to do so. In this case we get a much better estimate of the effect of packing than in either of our previous upper or lower bound models. The quick calculation we have made suggests that it would be worthwhile to try to repeat the calculation in three dimensions.

Try to draw Figure 2.4(b) in three dimensions and calculate the vertical distance between the centers of two packed spheres.

Alternatively, can you GUESS what the answer will be from the two-dimensional figure?

LEARNING TO LEARN

Were the two tasks instructive? How have you benefited from thinking about this problem? What have you learned about 'learning to learn' from this exercise?

Compile a list of points about HOW rather than WHAT you have learned?

Stop and Think

The 'ping pong' problem is one we have used with a large variety of audiences, from junior high school students to professional engineers. It originated (as far as we know) with a professor at MIT and came to us by way of Billy Koen. The problem is disarmingly simple but it helps develop several important skills and illustrates a number of sophisticated modeling concepts:

- it develops the courage to make 'back of the envelope' calculations and the wisdom to recognize when they are appropriate;

- it also develops the ability to match model resolution with available resources and to be critical of the resolution of a model;

- it encourages a pragmatic awareness of assumptions and the trade-offs between assumptions;

- it illustrates the power of symbolic representations.

There are several points to be made about the way in which the problem was presented to you, and the tasks you were asked to perform:

We deliberately provided you with progressively increasing resources. This should have reinforced the points we wanted to make about constraints, assumptions and resolution. We tried to create an environment in which you would perceive these points BEFORE we discussed them.

In a classroom we would have:

- asked each pair of students to describe the method they used to solve the problem in the five minute exercise;

- recorded the answers and a brief description of the methods on the board;

- asked questions such as 'Which is the best answer?' and 'Why do you prefer one method to another?' or 'What are the strengths of the different methods? And the weaknesses?'

Topics such as resolution can then be discussed, pointing out that with more resources (time and people) the answer is better. These discussions would highlight the points we have made in this chapter about resolution. The classroom experience would show up how some students used the same procedure both times but refined their measurements, while others changed methods when more resources were available.

Notice that either way (in this book or in the classroom) our approach is to ask YOU to DO something - be it to write, list, construct or discuss. This forces you to commit yourself to a position before you read or hear what we or the teacher has to say. If you have developed ideas that are similar to ours, you will have learned them far more effectively than if you had read them without thinking about them ahead of time. If your ideas are different, we hope you will either defend them (it could be frustrating trying to argue with an intangible author!) or think about why our ideas are better than yours.

Asking you to work in pairs also serves a purpose. You are more likely to think about the problem if you are interacting with another person, and the two of you are likely to both come away from that interaction better prepared than if you had worked on your own.
In this problem there was also a motive in asking you to work alone for the sixty second exercise and then cooperate on the five minute exercise. We wanted to expose you to the need to communicate, to talk explicitly about your model even if you did not realize at the time that you were building a model.

> Did you struggle with the assignments? Were you perplexed and unsure of yourself?

We deliberately tried to create an atmosphere in which you were likely to struggle and be perplexed. Struggling is a precursor of learning and being temporarily perplexed is a natural phase in problem solving. You must learn to struggle with that uncomfortable feeling of taking risks in exploring the unknown.

You must also learn to look back on that struggle and evaluate what you have done. George Polya, the famous mathematician, wrote a wonderful little book on problem solving entitled 'How to solve it: A new aspect of mathematical method.' In it he recommended 'looking back' as a heuristic that should come after other heuristics such as 'understanding the problem', 'devising a plan', and 'carrying out the plan'. 'Looking back' is a good heuristic for problem-solving, but it is a VITAL heuristic for learning.

> Finally, have you noticed how we have kept on asking you questions all over the place?

This too has been deliberate. We have been trying to encourage you to think about the problem, and also to think about how you are thinking about the problem!

The type of questions we asked are questions you should eventually be asking yourself.

****

The goal of this book was to create a classroom environment in which the students do something (such as formulate a model for a problem) and then discuss it, revise it, discuss it, etc. We attempted to do this by a "punctuated dialogue" in which we set a task for the reader and then ask the reader to **STOP AND THINK**, to formulate a model. We then ask questions of the reader or provide some information and ask for a revision or a reconsideration. When we want the reader to pause and reflect, we use the "mini-logo" **Stop and Think**.

All ten chapters in the book were pilot tested in two classes. Discussion among students proved to be very effective; therefore, we encourage readers to work with at least one other person on the book.

Conclusion

Punctuated dialogue, whether in the classroom or in a book, is an exciting way to get students involved in problem formulation and modeling. The modification of our approach in the classroom, developed over a period of more than fifteen years, prompted us to write a book so others could join in the excitement that we are enjoying.

The adoption of active teaching-and-learning strategies and an explicit emphasis on students constructing models of complex phenomena and building explicit representations of problems they were trying to solve both had handsome payoffs. Students constructing models using the iterative "progressive refinement" heuristic led to successful and satisfied students. Requiring students to discuss, compare and contrast their representations of problems led to a deeper understanding in general and to an appreciation of other students differences, strengths, and weaknesses.

References

1. Johnson, D.W., Johnson, R.T., and Smith, K.A. Cooperative learning: An active learning strategy for the college classroom, Unpublished manuscript, University of Minnesota, 1988.

2. McKeachie, W.J., Pintrich, P.R., Lin, Y-G. & Smith, D.A.F. Teaching and learning in the college classroom: A review of the research literature. National Center for Research to Improve Postsecondary Teaching and Learning, University of Michigan, 1986.

3. Smith, K.A., Johnson, D.W. and Johnson, R.T., Structuring learning goals to meet the goals of engineering education. Engineering Education, 1981, 221-226.

4. Smith, K.A. Educational engineering: Heuristics for improving learning effectiveness and efficiency. Engineering Education, 1987, 77, 274-279. Reprinted in The International Journal of Applied Engineering Education.

5. Smith, K.A. The nature and development of engineering expertise. European Journal of Engineering Education, 1988, 13 (3), 317-330.

6. Mackworth, N.H. Originality. American Psychologist, 1965, 20, 51-66.

7. Smith, K.A., Starfield, A.M. and Macneal, R., Constructing knowledge bases: A methodology for learning how to synthesize. In L. P. Grayson and J. M. Biedenbach (Eds.), Proceedings Fifteenth Annual Frontiers in Education Conference, Golden, CO, Washington: IEEE/ASEE, 1985, 374-382.

8. Smith, K.A., Wassyng, A. and Starfield, A.M., Development of a systematic problem solving course: An alternative to the use of case studies. In L.P. Grayson and J.M. Biedenbach (Eds.), Proceedings Thirteenth Annual Frontiers in Education Conference, Worcester, MA, Washington: IEEE/ASEE, 1983, 42-46.

9. Starfield, A.M., Butala, K.L., England, M.M. and Smith, K.A., Mastering engineering concepts by building an expert system. Engineering Education, 1983, 104-107.

10. Koen, B.V. Discussion of the method. Unpublished manuscript, University of Texas, 1985.

11. Starfield, A.M.; Smith, K.A.; and Bleloch, A.L., How to model it: Problem solving for the computer age, 1988, In Press.

KARL A. SMITH

Karl A. Smith is an Associate Professor of Metallurgical Engineering at the University of Minnesota. His Bachelors and Masters degrees are in Metallurgical Engineering from Michigan Technological University and his Ph.D. is in Educational Psychology from the University of Minnesota. He has served the Educational Research and Methods Division as a member of the Board of Directors, Chairman of the Nominating Committee, Vice Chairman for Programs, and most recently as Vice Chairman for Publications. He was co-chairman of the Technical Program at the 1985 Frontiers in Education Conference. Karl contributes regularly to ASEE Conferences and is a frequent writer and speaker on the use of active learning strategies including cooperative learning and structured controversy, on knowledge representation and expert system building, and on the use of personal computers as instructional tools.

ANTHONY M. STARFIELD

Anthony M. Starfield is a Professor in Civil and Mineral Engineering and Professor in Ecology at the University of Minnesota. He is also part-time incumbent of the Chair of Applied Mathematics at the University of the Witwatersrand, Johannesburg. He has over 20 years experience in teaching applied mathematics, computing and modeling. He is co-author with A.L. Bleloch of Building models for conservation and wildlife management, Macmillan, 1986; co-author (with J.C. Jaeger) of An introduction to applied mathematics, Oxford University Press, 1974; and co-author (with S.L. Crouch) of Boundary element methods in solid mechanics, George Allen & Unwin, 1981.