

Moral and mathematical lessons from a Rubik cube

Apart from humiliating a whole generation of adults in front of their more adept offspring, Rubik's brilliant invention has sparked off some new thinking about how people solve problems. Here British cubemaster David Singmaster draws together some of his reflections on the cube phenomenon

David Singmaster Problem solving, like other mental activities such as creativity, is hard to define but generally easy to recognise. Psychologists argue whether problem solving is the same as or just part of thinking or learning, and they also debate over "What is a problem?"

Without getting into these deep and murky waters, let me say that problem solving uses thinking, but that solving a genuine problem usually requires some creative act, ranging from the flash of insight that illuminates the entire problem to the glimmer of light that shows a direction to pursue (and which may be a will o' the wisp). Without creativity, problems reduce to exercises, though this is a gradual transition and one cannot draw any firm line between problems and exercises. To me, the creative aspect of problem solving is beyond pure thought, though again there is a continuous transition involved.

Unlike thinking, which has been discussed by philosophers and systematised by logicians ever since the Ancient Greeks, problem solving has only recently become an object of study. The only relevant ancient work is *The Method* of Archimedes which describes his mechanical analogy for finding areas and volumes. The earliest general discussions that I know of are in the works of 17th century philosopher, René Descartes, *Rules for the Direction of the Mind* and *Discours de la Méthode*. In the late 19th century, Henri Poincaré and Jacques-Solomon Hadamard wrote on mathematical creation, and more recently George Pólya and Gabör Szegő used problems as a teaching method and Pólya produced his famous *How To Solve It* (Princeton UP, 1971). But since the 1950s, there have been several hundred articles and books, generally written by psychologists, but also from mathematicians, engineers, advertising men and so on. Several dozen mathematics texts have appeared that use problems as the principal teaching technique.

It is natural for mathematicians to be the first to study problem solving but it surprises me that other disciplines have been so slow. Problem solving is an important part of the work of researchers of all descriptions, as well as medical diagnosticians, archaeologists, translators, historians, detectives, accountants, designers and

taxonomists, but few of these fields have contributed to the problem-solving literature (excepting detective fiction). Perhaps the practitioners have just considered problem-solving techniques as "common sense".

The current interest in problem solving arises from the recognition that:


- problem solving is a common everyday phenomenon
- problem solving can be taught and learned (and has not been in the past)
- problem-solving skills are applicable in most fields
- problem solving can be fun.

One pet theory of mine is that the decline in the study of Latin has accentuated this interest. People often used to claim that studying Latin developed the power of logical thought and the ability to solve problems, thereby equipping students to run countries. I believe this is a much inflated claim, but there may be some truth to it. Now that the study of Latin has declined, interest has focused on how problem solving is taught in mathematics classes. The "new math" movement, which emphasises understanding and uses more non-routine problems, has provided the main stimulus for this.

In many ways, problem solving is just common sense but we know how rare that is! The ideas of problem solving appear to be well known or obvious, but most of us will be able to recall situations where we spent hours or days on a problem whose solution was "obvious" once we found it. The systematisation of common sense is generally a convenience and has sometimes even provided a breakthrough in a subject. I find that when I am stuck on a problem, a simple contemplation of different problem-solving techniques can lead me to a new approach. Teachers of problem solving believe that problem solving can be taught, that it helps people in solving problems and that this is a good thing.

We cannot really discuss problem solving without an example, so I am going to discuss a rather large problem, namely Rubik's cube. The discussion presented has the benefit of hindsight: although each idea appears obvious and in an obvious sequence, most were not obvious at the time they were discovered, the steps were not considered in this order and the overall result does not seem obvious even now.

We start by picking up and examining the cube. We



Cube aficionado David Singmaster with some of the many varieties of cube that have hit the British market

London Express Service

read the instructions (if any) and ask friends what it does. At this first stage, it is essential not to go too fast, but now most people have already done so; they have already jumbled the cube and thereby discovered the basic mathematical problem: how do you get back to where you started?

Let us assume that we have been more careful and not jumbled before we looked. When we look, we see that the cube apparently consists of a $3 \times 3 \times 3$ (which we abbreviate as 3^3) array of 27 pieces. These are small cubes, called cubies or cubelets, with their outside faces (called facelets) coloured. The original state had the 3^3 facelets of a single colour on each face, with six different colours for the six cube faces. It takes some people a time to realise that the inside facelets are never seen, to see that all the faces can turn and that there are several distinct types of pieces.

There are 8 "corner" pieces, 12 "edge" pieces, 6 "centre" pieces (which are the centres of the cube faces and might be called "face" pieces) and a central piece that is never seen. These pieces move about as the faces turn, but we soon notice that the corner pieces move only to corner places, and so on. The types remain separate and are not intermixed. Further, the centres only turn in place and are never permuted. I have given the above observations explicitly because they are the necessary foundations for later investigation and many people have trouble observing them, even though others find them obvious.


The right notation simplifies the problem

Whether the cube has already been jumbled or has been carefully preserved, we will realise by now that one way to restore the cube is "just" to reverse the moves we made. This requires some means of recording moves. Mathematicians are used to doing this, but notation tends to frighten non-scientists. However, notation is simply a way of naming things and every field has its own specialised words and jargon, so the basic idea is really familiar to most people.

I devised the notation now most commonly used about December 1978. John Conway, a mathematician at Cambridge University, had already developed a similar but more poetic notation but I did not see this until the summer of 1979. The basis of my notation is the observation that the centre pieces are not permuted—they only turn in place. Thus we can always tell what colour a face should be by looking at its centre, regardless of how mixed up the cube is. It is natural to identify the faces by their centre colours but this soon runs into complications. So the six faces are labelled "Front", "Back", "Right", "Left", "Up", and "Down" and abbreviated to their initial letters.

Now we label all the pieces and positions of the cube. For example, the corner at the intersection of the U , R , F faces is labelled URF (or RFU or FUR), the edge at the intersection of the U and F faces is labelled UF (or FU) (Figures 1 and 2). These labels denote the shown positions and the pieces that were originally present but which wander away into other positions as the faces are turned.

We also use face names to denote face turns. R means a 90° clockwise turn of the R face as in Figure 3. (Clockwise is the most natural way for a right-handed person to turn the R face.) R^2 denotes RR which is a 180° turn and R^3 denotes RRR , also written R^{-1} or R' , denoting a 270° turn clockwise or a 90° turn anticlockwise. $R^4 = RRRR$ is the same as the identity operation, I , of doing nothing. Clockwise can be confusing on faces that are away from you—I always treat it as clockwise when looking straight at the face. RU now denotes the process of first turning the R face, then the U face (Figure 4).



Giant announcement for March 1982 cube championship in Budapest, Hungary

PopperFoto

The final aspect of notation is the question of representing the effect of a sequence of moves. For example, consider the cube after RU has been executed (Figure 4). This diagram does record the effect of RU on the pieces, but it is hard to see what is happening. Let us consider a single piece such as the edge FR . This piece is carried to the UF position, which we denote by $FR \rightarrow UF$. (Here the relative orders are important. The notation $FR \rightarrow UF$ means that the F facelet of the piece in the FR position is carried to U facelet of the UF position. This is the same as $RF \rightarrow FU$, but is *not* the same as $FR \rightarrow FU$.) By following along, we see that: $FR \rightarrow UF \rightarrow UL \rightarrow UB \rightarrow UR \rightarrow BR \rightarrow DR \rightarrow FR \rightarrow UF \rightarrow \dots$, as shown in Figure 5. As there are only a finite number of edge pieces and the process RU can take only one edge piece into a new "edge piece" position, this sequence must cycle back to the beginning and we denote this "7-cycle" as $(FR, UF, UL, UB, UR, BR, DR)$. The corners go through the sequence: $BRU \rightarrow DRB \rightarrow FRD \rightarrow UFL \rightarrow ULB \rightarrow UBR \rightarrow BDR \rightarrow \dots$, as shown in Figure 6, where the positions repeat after five steps, but with an anticlockwise twist, so we write this "twisted" 5-cycle as $(BRU, DRB, FRD, UFL, ULB)^-$. There is another corner affected that remains in place but twists clockwise, giving $URF \rightarrow RFU \rightarrow FUR \rightarrow URF \dots$, which we write as $(URF)^+$. (The choice of clockwise for $+$ is arbitrary and probably due to thinking of the turn R as clockwise.) We call this representation (a modification of the standard cycle representation of a permutation) the cycle representation of the effect of RU on the cube. If we apply RU to a jumbled cube, the notation $FR \rightarrow UF$ means that the piece at the position FR is carried to the position UF .

We have spent some time on the preliminaries, particularly the notation, because notation is essential for communication and because it is often neglected as a stage in attacking a problem. Of course a suitable notation rarely appears instantly and clearly at the start of a problem. It is only once the problem is well under way or even completed that one fully knows what the notation is required to do. Notations are generally taught long before their significance is apparent.

We are now (at last!) ready to start the actual solving of the problem, and will use several problem-solving techniques in no particular order. We have already noted that there are three types of piece: corner, edge and centre. Because the centres have no visible changes on a standard cube, we ignore them. (Here we are eliminating irrelevant aspects of the problem.) We are now left with two types of piece and we can split our problem into two parts: restore the edges and restore the corners. There may be some interaction between the edges and corners that prevents this split but it is a reasonable split to start with. We also observe that a piece can return to its position but in a different orientation; for example, RU twists the FUR piece. So we can try to split the restoration into another two parts: get the pieces in their right positions and get the pieces correctly oriented. There are now four

such problems which can be assembled into a solution in any order, provided we can solve a subproblem in such a way as not to disturb the solution of the previous sub-problems.

As a general strategy, this is fine, but it is not yet broken down far enough to be useful. If we begin thinking about the details of the cube, we can try to restore one piece at a time. However, there are 20 pieces (eight corners and 12 edges) to restore and a piece may be in many different positions with respect to its home position, so such a method will involve solving many subproblems (239 to be exact) and this is not a very human strategy, though it is perhaps the easiest for a computer.

The order in which the 20 pieces can be restored is almost completely arbitrary, so that each solver of the cube can devise his own strategy. (This is one of the features that elevates the cube beyond the ordinary puzzle.) Most people follow the recommendation in the instructions and try to restore one face first, then follow with the adjacent layer and the remaining face. In this way one can simplify the number of subproblems considerably and one can group them into about eight stages, each of which has about six sub-problems, and this becomes humanly comprehensible. (Psychologists believe that the human mind remembers information in batches or chunks of at most about eight items.)

These general considerations of strategy should interact with our exploration of the cube, though many would-be solvers will not start thinking about strategy for some time. The strategic considerations tell us that we should try to find processes that have simple effects. In particular, can we move one piece into a desired position while disturbing only a few others? Some such processes can be found by trial and error or by simply paying attention

to only a few pieces.

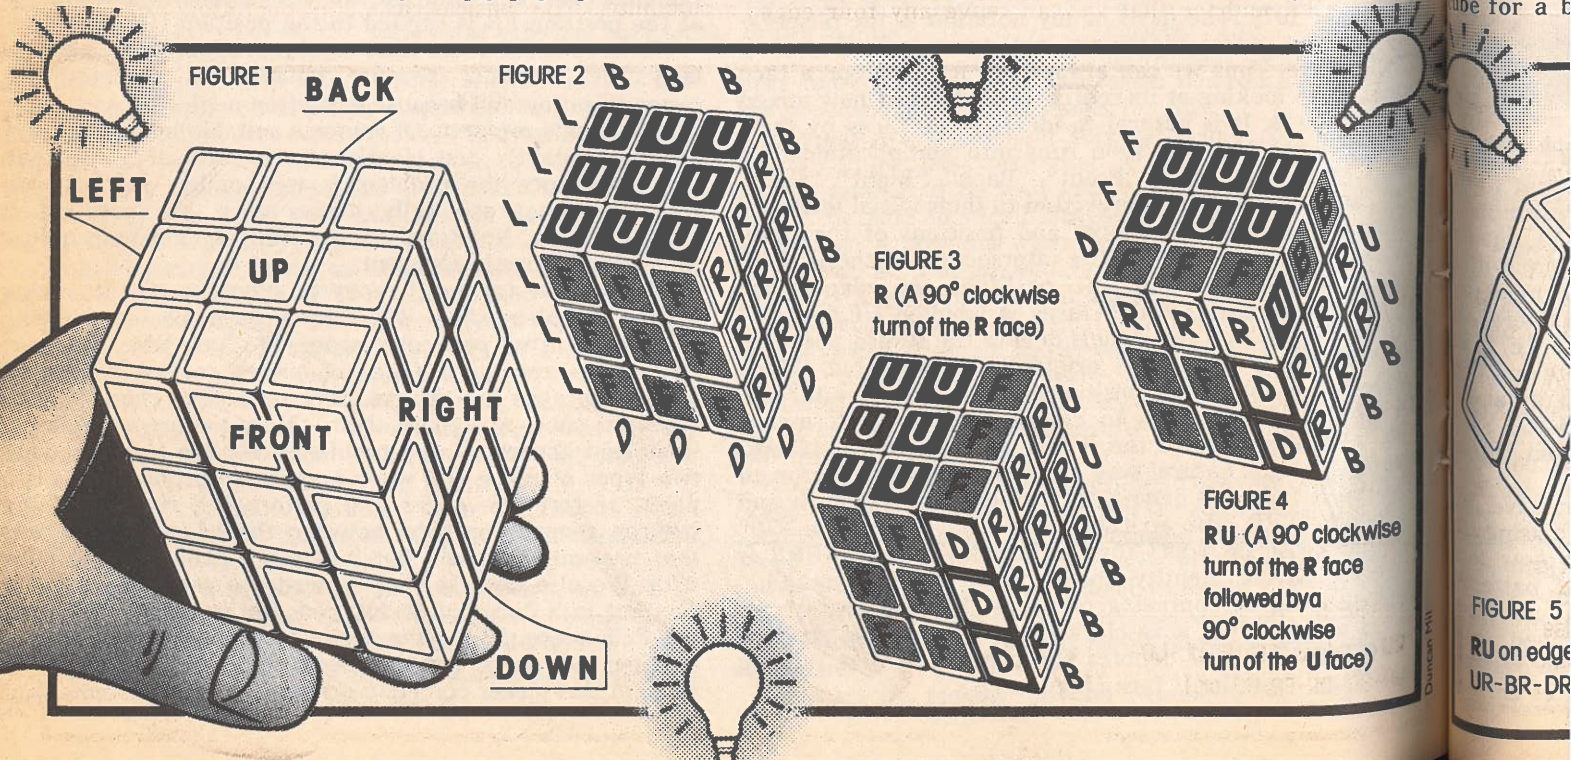
I found it easier to systematise my trial-and-error searching by restricting the possible moves, thus reducing a wider problem to a subproblem. The wider problem is to determine just what positions of the cube can be achieved by turning the faces. The associated subproblem is to determine what positions are achievable by turning, say, just two faces F and R . The set of such positions is denoted $\langle F, R \rangle$. The set of all positions of the cube is an example of the mathematical structure called a group; $\langle F, R \rangle$ is the subgroup generated by F and R .

Explorations among the subgroups

There are several such subgroups that one can easily explore: for example, $\langle F \rangle$ has only four elements and $\langle F^2 \rangle$ has only two. John Conway suggested the following subgroup which has many elegant patterns. A slice is a move like $F_s = FB'$, which can be viewed as the movement of the middle "slice" layer between the F and B faces (but then the relation of our coordinate system to the face centres is lost and one must use a "spatial" coordinate system which requires a bit more care to use). The subgroup of positions generated by the three slice moves, FB' , UD' and RL' , has only 768 patterns and one can soon find a simple way to return to the pristine cube from any one of these positions in at most seven slice moves.

The group generated by the "slice-squared" moves like $F_s^2 = F^2B^2$ has only eight elements. Playing with these groups does not lead toward our solution except that it develops familiarity with the cube and with the notation. It also yields a diversionary problem of continuing interest: what aesthetically pleasing patterns can be achieved?

If we now return to $\langle F, R \rangle$ and find it too complex (which it is), our experience with the slice moves suggests we consider $\langle F^2, R^2 \rangle$. We find this subgroup has precisely 12 patterns and one of them $F^2R^2F^2R^2F^2R^2 = (F^2R^2)^3$ has the particularly simple effect of exchanging the positions of two pairs of edge pieces (UF, DF) (UR, DR), as shown in Figure 7. In fact, as I shall discuss later, we cannot get a simpler effect. Exploration of different types of subgroups with one slice move will yield $U^2R_s^2D^2R_s^2 = (UF, UB)$ (DF, DB) and $F^2R_sUR_s = (UF, DF, UB)$ where R_s is the inverse of $R_s = RL'$, namely $R'L$. Note that the last 3-cycle is equal



and-error reducing problem is a problem turning situations is a group;
 Any one of these moves suffices to give lots of simple changes by means of a process that in group theory is called conjugation, though it is really a general problem-solving technique that has no general name. For example, how do I know how to wash a dish at the kitchen sink, how do I wash a dish at the table? Simple—I take the dish over to the sink and wash it, then bring it back to the table. On the cube, I know how to exchange two pairs of pieces in two particular positions by $(F^2R^2)^3$, so how do I exchange two other pairs of pieces? Simple—I take the other two pairs to suitable positions, then exchange them, and then put them back where I found them. If P is the process that brings the pieces into the working positions, then we must use P^{-1} to put them back again, so that our total operation is $P(F^2R^2)^3P^{-1}$. At first, finding P seems to be just as difficult as finding the two exchanges as P also moves four pieces. However, we can let P affect these other pieces, because $(F^2R^2)^3$ does not affect any other pieces and so P^{-1} will also put them back correctly. Hence a conjugate of $(F^2R^2)^3$, such as $P(F^2R^2)^3P^{-1}$ is also a pair of exchanges and it acts only on the pieces that are brought to the four working locations by P . My friend and colleague Sandy Frey has correctly recognised this as a fundamental process for the cube, and for other mathematical groups, and he calls it the *Principle of Partial Inverses*; a piece moved by P is restored by P^{-1} , provided nothing else has moved it in between.

Change to a problem you can solve

Conjugation is generally taught in algebra or group theory as a fairly abstract process. I certainly never pictured it as being the same operation acting on different pieces until I actually saw it happening on the cube. While writing this article, I have realised that it is actually a very general problem-solving technique, of which there are many other instances in mathematics and science: If you cannot do a problem, transform it to a problem that you can do and then transform the answer back to reach the solution to the original problem.

Conjugation is a powerful technique on the cube and in general, but how shall we apply it? We contemplate the cube for a bit and see that we can move any four edges

into into the four working locations of $(F^2R^2)^3$, and we can put the edges in either way round.

In this way we can put one edge at a time in place while we exchange another pair of edges. This procedure may break down if we get to a single pair of edges, but we will see that this can happen only if a pair of corners has also to be exchanged. The fact that we can put the edges in the working locations either way around allows us to exchange a pair twice, but with one edge flipped between exchanges, and this leads to all edges being in place but one pair being flipped. With this, we can easily flip any even number of edges, which we will see is all that is possible.

In principle, then, we can restore the edges of the cube without affecting the corners! This is quite a lot to achieve with a single process. It certainly demonstrates the power of conjugation and provides a definite guide for further puzzles of this type. When I first acquired a cube, Conway and Roger Penrose, from Oxford University, had already solved theirs and remarked that it took them some time to realise the power of conjugation. It certainly was the last point that I understood; I remember lying awake thinking about it, seeing that I could move any four edges into the working locations and realising that this completed the general method for restoring the cube to its original state.

What remains is to solve the problem of restoring the corners of the cube to their rightful positions. We return to our subgroup $\langle F, R \rangle$ and look at another common construction of group theory. The commutator of F and R is $FRF'R'$ which I denote $[F, R]$. First, let us see why this is called a commutator. The inverse of FR is $(FR)^{-1} = R'F'$ so that $FRR'F' = I$. Thus $FRF'R'$ is not normally the identity; it is the identity, if and only if $FR = RF$, that is, F and R commute. The basic arithmetic operations of addition




FIGURE 5
 RU on edges (FR-UF-UL-UB-UR-BR-DR-FR-UF-etc)

FIGURE 6
 RU on corners (BRU-DRB-FRD-UFL-ULB-UBR-BDR-etc)

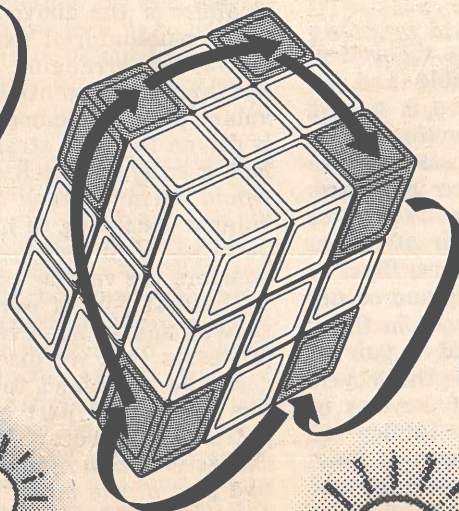



FIGURE 7
 $(F^2R^2)^3 = (UF,DF)(UR,DR)$



Duncan Mill



and multiplication are commutative: $a + b = b + a$, and $ab = ba$. Though subtraction ($a - b \neq b - a$) and division ($a/b \neq b/a$) are examples of non-commutative operations, the following may be more appropriate. When we get up in the morning, we put on our socks, then our shoes. To go to bed, we take off our shoes, then our socks to get back to the original state of bare feet. That is, we must reverse each step and the order of the steps. Trying to take off our socks first will cause a terrible mess. On the other hand (or foot), consider putting on the right shoe (R) and then the left one (L). The shoes can now be taken off in either order, or we can see that the order of putting them on makes no difference, that is, $RL = LR$, or R and L commute. R and L also commute on the cube because they act on different pieces, but F and R do not commute, so the commutator $[F, R]$ creates a terrible mess (Figure 8).



Or is it such a mess? F and R produce cycles of corners and edges which overlap in the three positions along the FR edge of the cube. The Principle of Partial Inverses shows that $FRF'R'$ does not affect a piece like FL because F does not carry it into the overlap region, hence it is unaffected by R, then restored by F' and again left unaffected by R' . From this we can see that the only pieces that can be affected by the commutator $[F, R]$ are those which are in the overlap or are moved into the overlap by F or R. As there is just one edge in the overlap region, the action on edges must be a 3-cycle.

The overlap has two corners, one of which moves to the other under F and under R, so we find that the commutator is a pair of twisted 2-cycles, affecting only four out of the six corners involved in the moves F and R. The same sort of reasoning about conjugation that we used for edges shows that we can use this process to restore all the corners, though we do disturb the edges. However, we can restore the corners without disturbing the edges, and once the corners are correct, we cannot find ourselves left with a single pair of edges to exchange (as we shall see later). In theory, this completes the constructive part of our problem; every possible pattern of the cube can be restored by these processes.

How to find new tricks

However, we have not yet demonstrated the theoretical assertions that certain patterns are impossible and we shall do this shortly. In practice, our solution is a long way from effective. But the notion of the commutator is very powerful for finding useful simple processes. If we can find a process that affects only one piece in a face, then the commutator of that process with that face gives a 3-cycle of a piece, or a pair of flips or a pair of corner twists, depending on whether the process moves, flips, or twists the piece. Noting that $[F, R]$ moves only one corner in the L face, and $[F, R]^2$ twists only one corner in the L face, we can obtain a 3-cycle of corners and a pair of corner twists. We also note that $[F, R]^3$ causes the 3-cycle of edges to vanish and leaves just a pair of 2-cycles of corners. These results greatly improve our ability to restore corners and we can now restore corners and twist corners without affecting edges. With these processes, our

edge processes and the technique of conjugation, one can actually solve all four of our basic subproblems independently of the others and it is possible to do them in any order, although it is difficult to orient pieces that are not yet in their correct position.

We can now describe the possible patterns of the cube precisely, but it is easiest to first contemplate a more general problem. Consider a cube which is disassembled into its pieces and imagine reassembling the pieces so that the outside facelets remain outside and the centres remain in the same relative positions. (This can actually be done because the corner and edge pieces can be removed, leaving the centres attached together, and one can put only corner pieces into corner locations, etc.) There are eight corner pieces and these can be rearranged or permuted in $8! = 8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1 = 40\,320$ ways. Each corner piece can be oriented in three ways, giving $3^8 = 6561$ orientations of corners. Similarly, we get $12! = 479\,001\,600$ and $2^{12} = 4096$ for the edges and a total of $8! \cdot 3^8 \cdot 12! \cdot 2^{12} = 519\,024\,039\,293\,878\,272\,000$ patterns which one could construct in this way. However, our conservation laws tell us that we can achieve just half of the total permutations of corners and edges, just half of the total flips of edges and just one-third of the total twists of corners, so we can achieve just one-twelfth of the above numbers, yielding $43\,252\,003\,274\,489\,856\,000$ possible patterns for the cube. A modern computer can count somewhere between 10^6 and 10^9 numbers per second, and there are about $3 \cdot 16 \times 10^7$ seconds in a year. Even at the faster



Minh Thai, 16-year-old winner of the Budapest championship

rate, it would take over 1300 years simply to count to the smallest of the above numbers. This is a nice example of the "combinatorial explosion": the fact that the number of ways of arranging n things grows very rapidly with n and becomes essentially infinite even for quite moderate values of n . To demonstrate the size of such numbers, one is driven to make strange comparisons; for instance, if we had a cube for each of the 4.325×10^{19} patterns, they would form a stack about 260 light years long; and a volume consisting of one cube for each of the 8.858×10^{23} patterns with marked centres would be about three-quarters the volume of the Moon.

The first part of our solution has been constructive. We showed how to achieve certain basic patterns by using the processes that we discovered and then we saw how they could be combined to achieve a large class of patterns. The theoretical part of the solution shows that no other patterns can occur. The theory can be described as a characterisation of the possible patterns or as the destructive part of the problem—it destroys or eliminates certain

possibilities. I
tom up an
possibilities, I
e known a
e construc
terns while
ossible patte
own intern
ructive and
cesses con

que comes
y to prove
ue may be
question.
In a sense,
at an impo
ection. Can
r results?
nary probl
tend our w
Can we n
a method
histlethwait
outh Bank
most 50 m
ounting sho
moves
21 right
be have be
restored in 1
enerated by
15 moves.
the latter re
for the comp
at one can
moves. The r
efficient way
remain know
Can we re
any genera
 $\times 3 \times 3$ Mag
ther polyhe
piece puzzle:
generally ea
more positio
on is an e
of the Unive
between the
corner and 1
perhaps the
tricate and
nd, to a les

...one can
independ-
m in any
t are not

...the cube
a more
assembled
pieces so
centres
actually

...ibilities. I also like to view these parts of the theory as
bottom up and top down. We start off with a set of *a priori*
possibilities, namely the set of constructible patterns, and
the known achievable state, namely the initial position.
The constructive theory builds up the set of achievable
patterns while the destructive theory reduces the set of
possible patterns. The problem solver pushes into the un-
known intermediate region, alternatively employing con-
structive and destructive processes. When these two
processes converge, the problem is completed.



...one can
one can
There
anged or
20 ways.
s, giving
we get
a total of
ns which
conserva-
f of the
lf of the
al twists
he above
sible pat-
nt some-
nd there
e faster

...ique comes into play: if you cannot prove it true, then
try to prove it false. In substantial problems, this tech-
nique may be used many times before one finally resolves
the question.

In a sense, we have now reached the end of our problem.
But an important aspect of problem solving is retro-
spection. Can we improve our results? Can we systematise
our results? Can we find simpler proofs? Are there diver-
sionary problems which could now be attacked? Can we
extend our work to similar problems?

Can we now quantify our qualitative results? There
is a method of solving the cube, due to Morwen
Thistlethwaite, my colleague at the Polytechnic of the
South Bank in London, that always restores the cube in
at most 50 moves (recently reduced from 52) and simple
counting shows that some positions require at least
38 moves (counting 180° turns as single moves)
or 21 right-angle turns. Several subgroups of the
cube have been completely analysed. The 2³ can always be
restored in 11 moves or 14 quarter turns. The subgroup
generated by the 180° turns on the 3³ can always be solved
in 15 moves. Both these results are the best possible and
the latter reduced Thistlethwaite's method by two moves.
For the complete cube, the results so far lead us to believe
that one can always restore the cube in at most about 22
moves. The method that restores any position in the most
efficient way is called God's algorithm and seems likely to
remain known only to God.

Can we relate the cube to other work? The cube has
many generalisations, such as the 2³, the 4³, the 3⁴, the
3³×3³ Magic Domino, and the basic idea extends to
other polyhedral shapes. The ideas also cover most sliding
piece puzzles. They require some basic work but are
generally easier than the cube, though some have many
more positions. A much more interesting type of rela-
tion is an elegant analogy, due to Solomon W. Golomb
of the University of Southern California at Los Angeles,
between the unachievable twists of a third turn of a cube
corner and the unobservable quarks of charge one-third.
Perhaps the most interesting uses of the cube are as an
intricate and interesting concrete example in group theory
and, to a lesser degree, in combinatorics. Almost all the




FIGURE 8
F,R = FRF'R' = (FLU, FUR) + (FRD, DRB) - (FU, FR, DR)

basic ideas of group theory occur naturally as one studies
the cube and a fair amount of advanced group theory
emerges. Already several scientific papers inspired by the
cube have been published: so has a textbook using the
cube to develop theory for students.

From the point of view of problem solving, I find the
cube has led me to formulate many problem-solving
techniques in more generality than previously and to
develop a general scheme for all problems of this sort.
The scheme is as follows:

1. Familiarise yourself with the problem
2. Find appropriate notation(s) for pieces, positions and moves. (In general, this may require assuming all pieces are distinct even when they are indistinguishable.)
3. Look for simple processes: especially pairs of 2-cycles or single 3-cycles, by using commutators and other subgroups
4. Once simple processes are found, investigate whether conjugation will allow the processes to be moved to other pieces
5. Look for parity or other conservation laws.

So far, this scheme has worked with every puzzle of this
type. The hardest problems are those where simple
processes are hard to find, usually because of some com-
plexity of the puzzle which makes it hard to see the
effects of moves, or those where the basic coordinate
system and hence a notation is hard to find. The scheme
yields only a theoretical solution, from which the number
of patterns can be obtained. Practical solutions require
finding more processes and developing a strategy, but
some puzzles become much easier because they have many
pieces the same.

So where does this essay leave us? I have tried to con-
vince you that problem solving is important and that it
can be taught and learned by means of guided problem
solving. In the example of Rubik's cube, we have looked at
a substantial problem that leads naturally to a great deal
of mathematics and uses many techniques of problem
solving. The cube leads to more mathematics and uses
more problem-solving techniques than any other recrea-
tional problem, and so is an excellent problem for teaching
and learning. □

Duncan Mill