THE FAILURE OF THE MIND'S EYE

We humans pride ourselves on being an intelligent species, but in at least one area of intelligence—spatial perception—we are evidently not as smart as we think. If you've taken an intelligence test or studied mechanical drawing, you've undoubtedly encountered problems in spatial perception. Given a perspective drawing of a geometric solid, imagine what the solid looks like in three dimensions. Some of these problems seem fairly simple and have been known for decades, but recent research at IBM shows that they have multiple solutions—solutions that even able mathematicians had failed to find.

Computer Vision

At IBM's Thomas Watson Research Laboratory in Yorktown Heights, New York, Michael Wesley was working with George Markowsky, who is now at the University of Maine, in the field of computer-aided design. They were teaching computers to "visualize" three-dimensional objects based on engineering drawings. Wesley and Markowsky developed a computer algorithm that, given a two-dimensional projection of the edges of an object, discovers all objects with that projection. In its current incarnation, the algorithm is restricted to objects whose edges are straight and whose faces are planar. Such objects—cubes, rectangular solids, prisms, pyramids, and so on—are known to the mathematical cognoscenti as polyhedrons. When Wesley and Markowsky gave the algorithm certain classic problems in spatial perception, to their astonishment it found multiple answers.

I went to visit Wesley, and we were joined by his colleague David Grossman, who has followed Wesley's work from the beginning. I returned home two hours later with a briefcase full of diagrams of polyhedrons that they had kindly sketched for me in order to elucidate Wesley and Markowsky's work. The account that follows is based on their sketches and their explanation of them. It took me many hours, and several false starts, before I could visualize all the shapes they were playing with—so do not blame yourself if your mental vision is less than 20/20.

Let's start with a simple problem in spatial perception. Suppose I tell you that the top view (a view is not necessarily what you see, but is a two-dimensional projection of edges), the front view and the side view of a polyhedron are all

> At IBM, research in computer-aided design has proved something unsettling about us: We're not especially skilled at spatial perception. You can see for yourself by trying the puzzles on the next page.

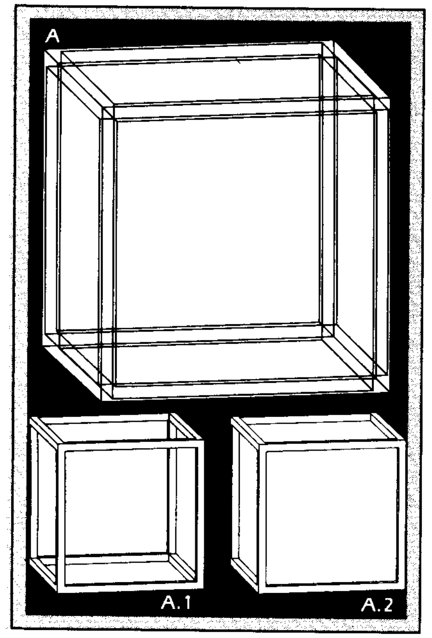

In diagram A the edges of an object are shown in space. What is the object? Well, it could be a tubular structure (A.1) or it could be a solid cube that has a tubular border (A.2). IBM researchers are training computers to solve problems like this one.

squares (diagram B) Can you figure out what the polyhedron is? (Those of you who are not familiar with engineering drawings should keep in mind an important drafting convention: Hidden edges—edges that are inside the object or on the other side—are represented by dotted lines unless they fall directly in back of solid lines, in which case they are not shown. Here there are no dotted lines, so at least this kind of chicanery need not be looked for.)

The answer, of course, is a cube (diagram B.1). But suppose I eliminate the side view, showing you only the top and front views (diagram C). Now what's the polyhedron?

Most of you will still say a cube, but that's only one of five possibilities. The other four are prisms (triangular solids), each of whose side view is a right triangle in one of four orientations (diagrams C.1, C.2, C.3 and C.4). Try to convince yourself that this is the case.

Now comes the fun part. Suppose the top view is a square with a horizontal line crossing it and the front view is a square with a diagonal (diagram D). What's the polyhedron?

If you came up with an answer at all, you probably visualized a rectangular solid joined to a prism. From the front the diagonal edge is visible, and from the top the horizontal ridge (namely, the top front edge of the rectangular solid) can be seen (diagram D.1). What few people realize is that nine other polyhedrons also have the same front and top views. For example, a new polyhedron (diagram D.2) can be formed by lopping the back off the polyhedron in diagram D.1. To be specific, imagine that a triangular slab is excised from the rectangular solid in D.1 by cutting it in half along the diagonal running from the top back edge to the bottom front edge. In other words, the rectangular solid is converted into a prism, the polyhedron as a whole becoming two linked prisms. Try to visualize in your mind's eye that a third polyhedron with the same front and top views can be formed by turning D.2 upside down (diagram D.3).

The Pharaoh's Chamber

Suppose the views are complicated by adding a second horizontal to the top view and a second diagonal to the front view (diagram E). Among spatial-perception groupies, this is known as the two-ramps problem. It was reputed to have 12 solutions until Wesley and Markowsky's algorithm found 107, five of which are shown in diagrams E.1, E.2, E.3, E.4 and E.5. "From a human point of view," Wesley said, "there can be all sorts of things that humans don't think of—cavities and odd juxtapositions that take a lot of puzzling to visualize."

Consider, now, top and front views that are squares within squares (diagram F). What polyhedrons are possible? The drafting convention of representing hidden

DIAGRAM COURTESY OF MICHAEL WESLEY, THOMAS WATSON RESEARCH LABORATORY, IBM

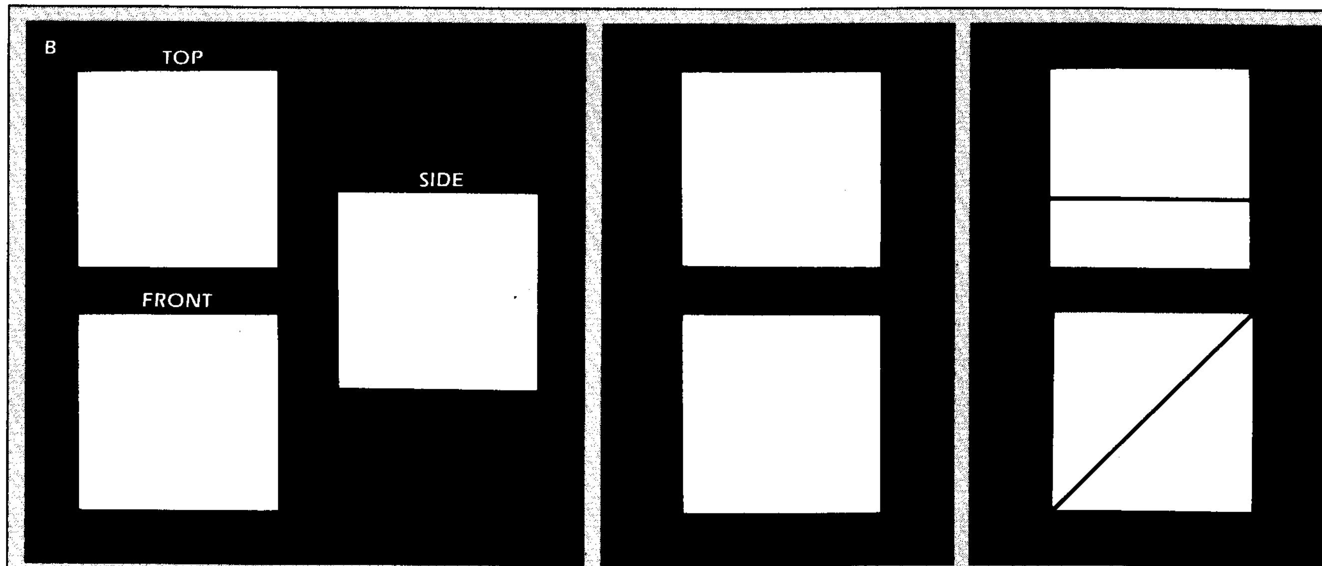

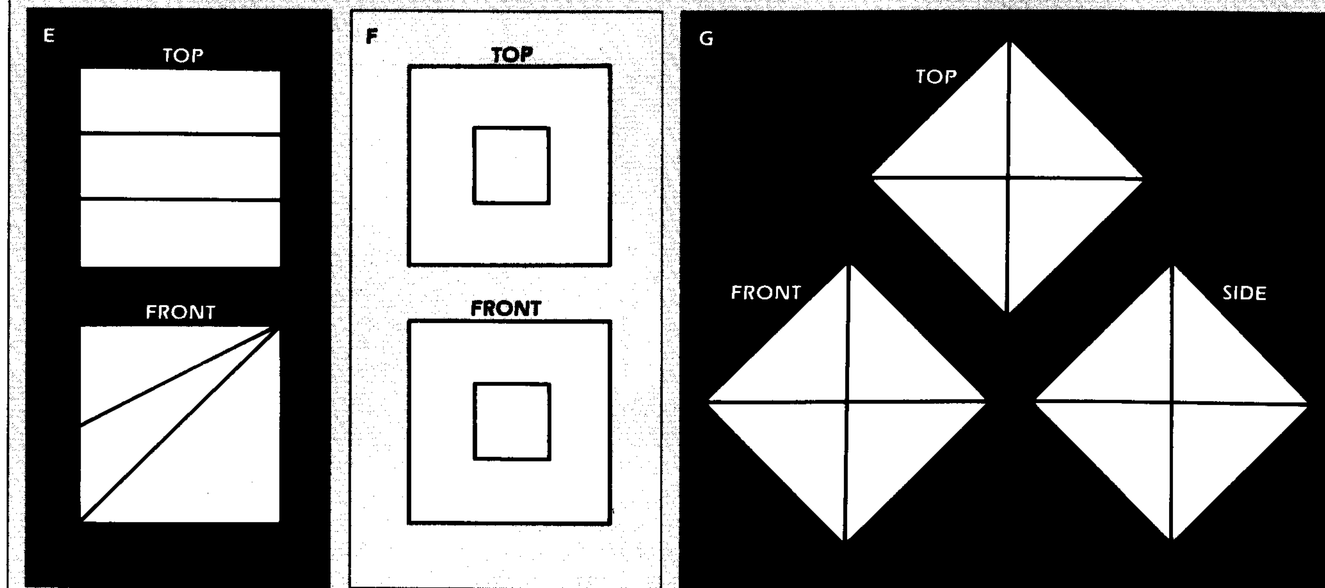

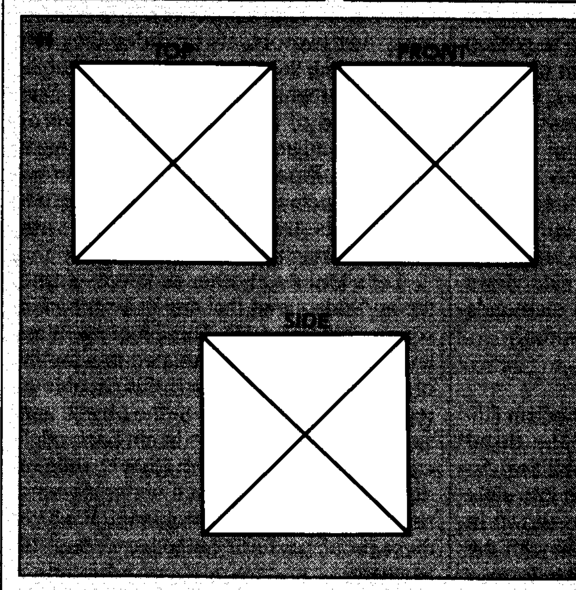

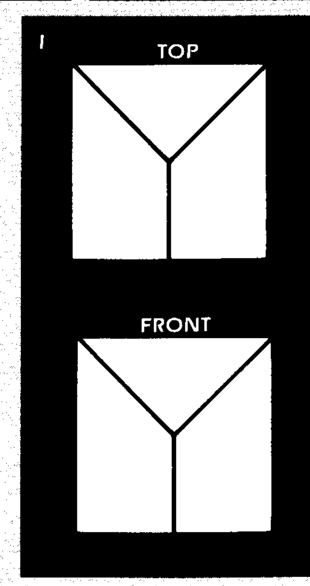

ALL DIAGRAMS BY MICHAEL WESLEY, REDRAWN BY NINA WALLACE

TEST YOUR SPATIAL PERCEPTION

Each of these diagrams shows the top and front view (and sometimes the side view) of an object (or objects) with straight edges and planar faces. You are now urged to try to determine what the objects are. The answers are shown on the next two pages and are explained in the text. The text also discusses how mathematicians and computers have fared on the problems.

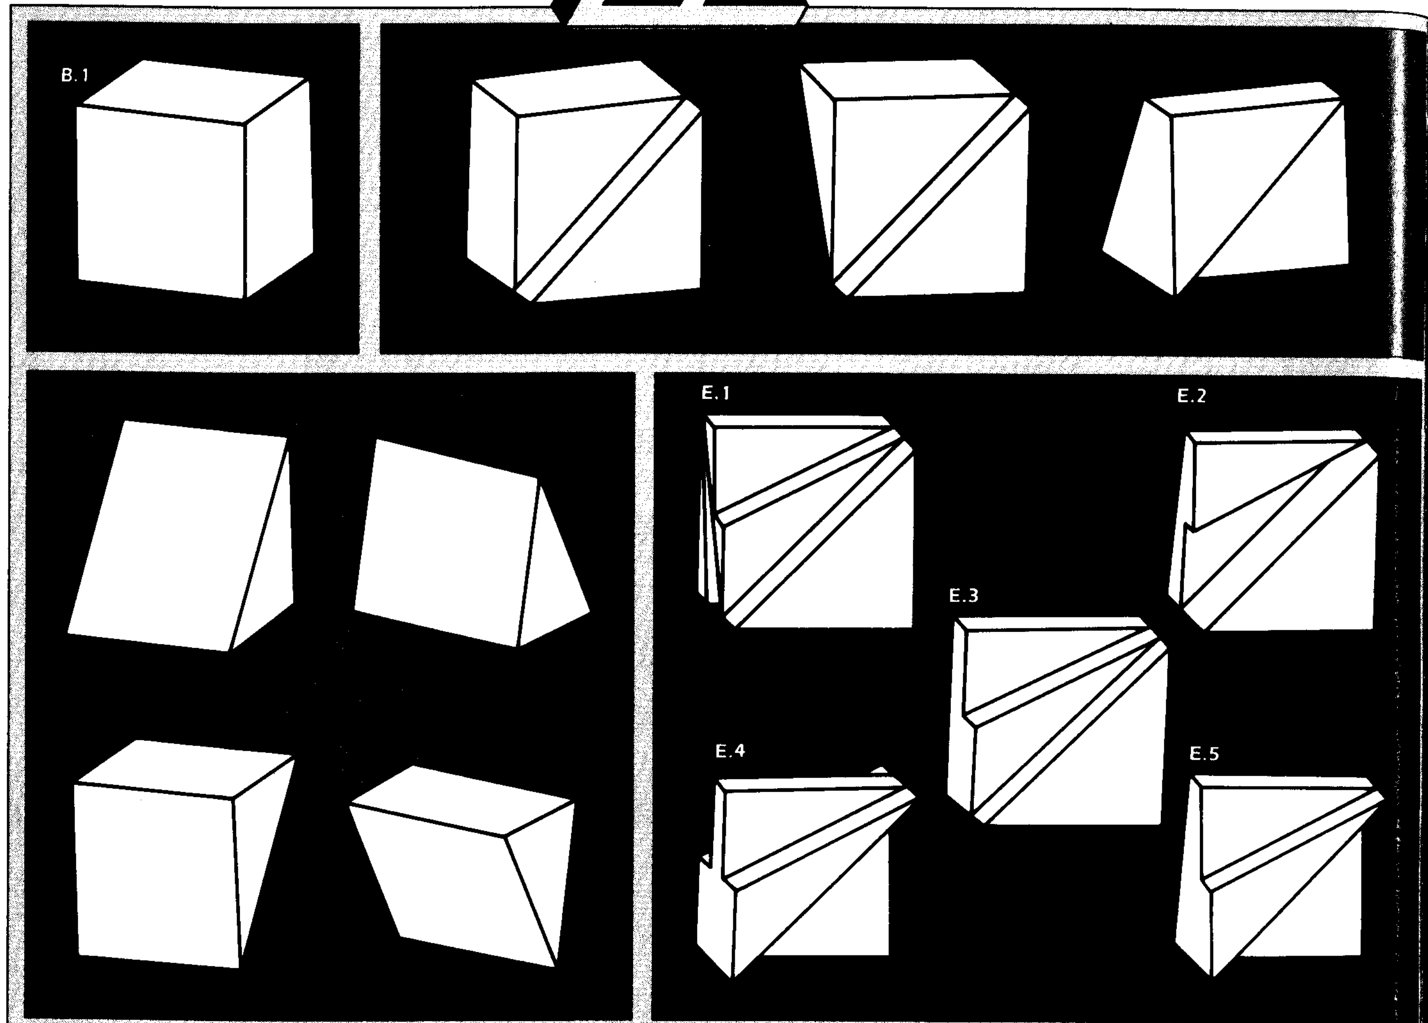

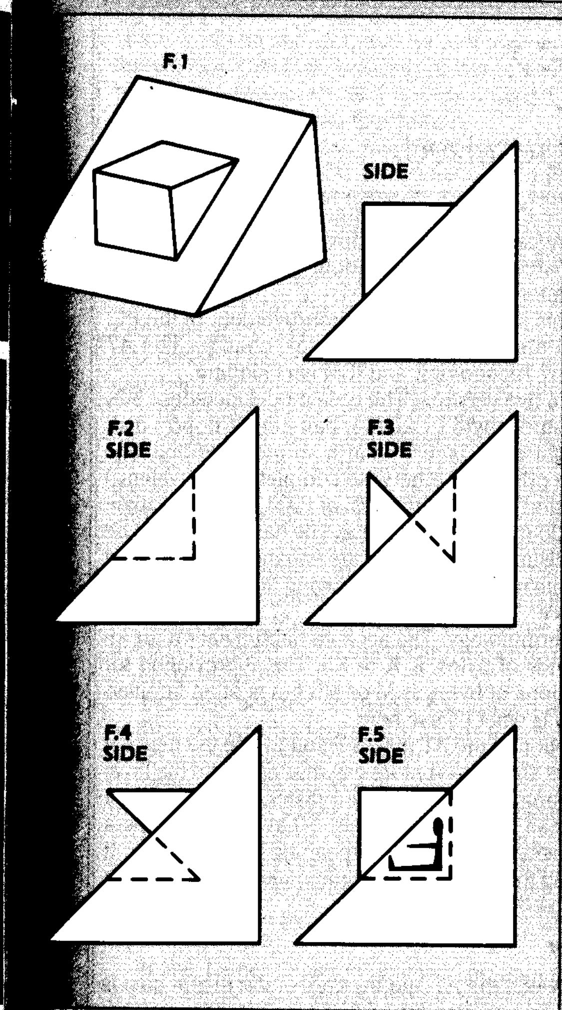

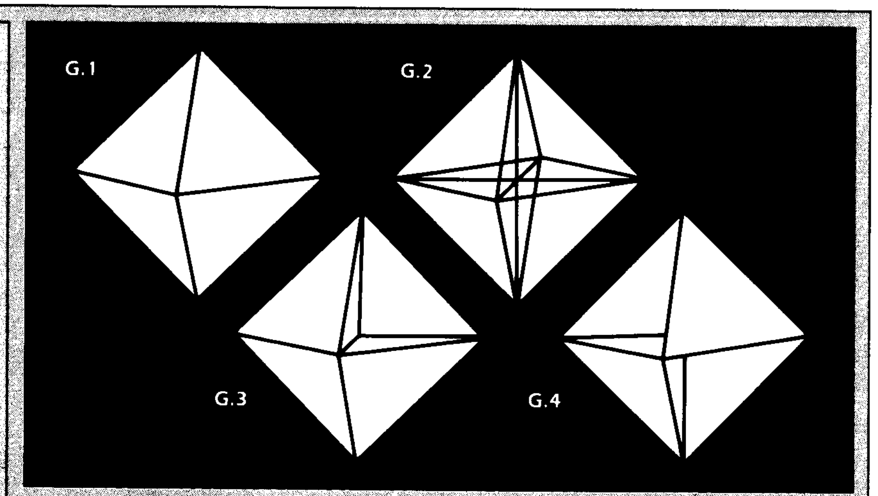

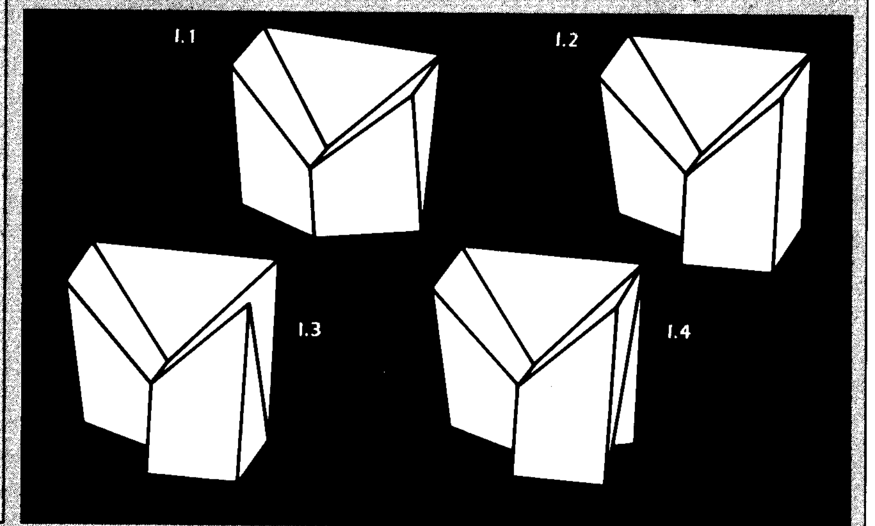

edges by dotted lines rules out a cube with a hollow cube at its center; the front and top views of such an object would be dotted squares within solid-line squares.

The most common answer—a right-triangular solid with a smaller right-triangular solid attached to it—is shown in diagram F.1. From the top you see a square and from the front you see a square, but looking from the side you see two right triangles of differing size. Another polyhedron (diagram F.2) has a small triangular hole instead of a small triangular wedge. Two other amusing possibilities, which most humans fail to think of, combine the idea of a hole and an appendage (diagrams F.3 and F.4) . But the best is yet to come.

"I remember," said Grossman, "a conversation that Mike and I had about this problem a long time ago. He told me there were five possibilities. But I could only see four. The fifth one I would not have found. Do you see what it is? It's like an old Egyptian Pyramid, with a hidden chamber for the pharaoh [diagram F.5]. It's got a bump [like the polyhedron in F.1]. But there's a cavity underneath it that you don't know is there. That's where the pharaoh sits."

At this point I detect a suspicion that many of you undoubtedly harbor. You suspect that the only reason we're getting multiple answers is because I'm cheating by showing you only two views, the top and the front. The one time so far that I've showed you the side view too, for the cube (diagram B), there was only one answer. Now let's test the hypothesis that three views—top, front and side—are needed to specify the object uniquely. Suppose I show top, front and side views that are plus signs inscribed in diamonds (diagram G). What's the polyhedron?

Faulty Hypothesis

The obvious answer is an octahedron (diagram G.1). But, in spite of the three views, there are 34 other possibilities. An octahedron can be decomposed into eight octants, the edges of which are shown in diagram G.2. Other polyhedrons with the same top, front and side views are formed by throwing away certain of these octants. The question is, how many can you throw away? You can discard any single octant; this accounts for eight solutions, one of which is in diagram G.3. You can throw away any two octants provided they are not adjacent; this gives rise to 16 possibilities, one of which is in diagram G.4. You can get rid of three octants, no two of which are adjacent; this contributes eight solutions. Finally, you can even throw away four octants, no two of which are adjacent; this gives rise to two more answers. So much for the hypothesis.

Let's look at another problem—front, top and side views that are *X*s inscribed in squares (diagram H). This problem, Wesley explained to me, is related to a family of problems. Think of each *X* as capable of representing four edges in the object, and so 12 edges are possible in all. Each edge can be exposed or hidden (that is, solid or dotted in the perspective views), giving rise to a family of 2^{12} problems. Each of these problems turns out to have about 10 answers.

"The weekend I was trying to find the solutions," Wesley said, "I started the computer running on a Friday night. Every time it found a solution, it was supposed to write it in a file. Marvelous, I thought when I came in Monday morning. Let's see what it found over the weekend. I looked in the solution file, and it was empty! It was a great letdown. I knew it had been running for two and a half days, so it just *had* to have found something. I realized then that the algorithm needed some improvement. It was spending much too long looking at things it could ignore." Wesley and Markowsky were able to improve the algorithm to the point where it can now efficiently find the polyhedrons.

Diagram I shows a drafting problem—the two *Y*s—that has been around for at least 30 years. The top and front views are *Y*s inscribed in squares. "I was given this problem," Wesley said, "as an engineering student in drawing classes at Cambridge—that would have been in 1957 or '58. I gave a talk about two years ago at the engineering department at Cambridge, and I showed them this diagram and explained how I had seen it as a student. A voice in the back said, 'We still get it!' "

"It's a very nice problem," added Grossman. "Unlike the others, the solution involves some planes at strange, dihedrally dipped angles. So it's not obvious. And when you start constructing it in your mind, most people have difficulty visualizing that those are planes, even when you show them where the edges are. It turns out that the problem doesn't have just one answer. And the fact that it has more than one answer wasn't realized, to my knowledge, until about ten years ago, when someone came up with a second solution. Smart people typically take two hours to find just one solution [diagram I.1]."

Seven Solutions

Wesley and Markowsky's algorithm discovered seven solutions. "I think that's pretty nifty," said Grossman, "since humans have a difficult time finding even one." Besides the one that people find, there are three solutions (diagrams I.2, I.3 and I.4), each having a mirror image.

"A neat thing about this problem," said Wesley, "which Markowsky discovered, is that the center point of the *Y* must be exactly in the middle of the square. In fact, if there is the slightest deviation, there are *no* solutions" because a plane gets a twist—and curved surfaces are not allowed in a polyhedron.

Lest you think that the multiplicity of answers to these problems stems, somehow, from the fact that top, front and side views don't contain enough information about an object, I have a surprise for you. Even if you know where all the edges are in three-dimensional space, the object may still not be uniquely defined. For example, diagram A shows all the edges of an object. What is it? Well, it could be a tubular structure (diagram A.1) or it could be a solid cube with a tubular border (diagram A.2). And this is nothing. Wesley has sketches of all sorts of horrendously complicated edge configurations, but I escaped before he had a chance to show them to me. ■