Helping Users Think in Three Dimensions: Steps Toward Incorporating Spatial Cognition in User Modelling

Michael Eisenberg*, **Ann Nishioka***, **M. E. Schreiner**†
* Department of Computer Science and Institute of Cognitive Science
† Department of Psychology and Institute of Cognitive Science
Campus Box 430
University of Colorado, Boulder CO USA 80309-0430
duck@cs.colorado.edu; (303-) 492-8091

ABSTRACT

Historically, efforts at user modelling in educational systems have tended to employ knowledge representations in which symbolic (or "linguistic") cognition is emphasized, and in which spatial/visual cognition is underrepresented. In this paper, we describe our progress in developing user models for an explicitly "spatial" educational application named *HyperGami*, in which students design (and construct, out of paper) an endless variety of three-dimensional polyhedra. This paper gives a brief description of the HyperGami system; discusses our observations (and experimental results) in understanding what makes certain polyhedral shapes difficult or easy to visualize; and describes the ideas through which we plan to augment HyperGami with user models that could eventually form the computational basis for "intelligent spatial critics."

Keywords:
Spatial cognition, user modelling, HyperGami, polyhedra.

INTRODUCTION

It has long been an ambition of software designers—particularly designers of educational computing systems—to incorporate user models in their applications. Indeed, it is arguable that the entire field of intelligent computer-aided instruction (ICAI) is based ultimately on the notion of creating computational models of students' knowledge, skills, or goals; these models may then be employed to diagnose errors or misconceptions, chart progress, or offer guidance or advice. Classically, such models have employed schemes in which students' knowledge is represented in symbolic structures—as a collection of productions [1], as a structured "lattice" of skills (expressed as production-like rules) [5], as a semantic network of facts (expressed as relationships between symbols) [7], or as a set of "issues" with which the student is familiar [6]. Perhaps not coincidentally, the domains to which these models are applied—the domains for which the ICAI systems are thus written—tend to anticipate the knowledge representation schemes to be employed. These systems consequently focus on domains in which expertise seems to be most easily expressed through natural language—domains such as game-playing [6], subtraction [5], algebraic manipulation [1], geographical knowledge [7], programming [1, 17], and so forth.

Important as these "linguistic" representations of knowledge are, there is an increasing body of evidence to suggest that spatial/visual cognition is an important factor in expertise in a variety of mathematical, scientific, and artistic domains. Gardner, for example, includes spatial intelligence as one of his "multiple intelligences" [14]; while Siemankowski and McKnight [25] present evidence that spatial cognitive abilities show a strong correlation with performance in college science curricula. Undoubtedly, in the lore of creative scientific and mathematical work, visual imagery and thinking are accorded a high importance: in a marvelous book, Miller [19] presents historical evidence indicating that such thinkers as Boltzmann, Einstein, and Poincaré thought of their own cognitive processes as strongly visual in nature; Hadamard [16], in a classic description of mathematical cognition, places a heavy emphasis on visual thinking (see also Fomenko [13] for a more recent expression of similar sentiments); while Ferguson [11] discusses the role of visual thinking in engineering and design.

This paper follows up on the work of these researchers, describing our progress in developing user models for a domain that is strongly spatiovisual in nature: namely, the construction of three-dimensional polyhedral solid shapes. Our motivation for these user models arises from efforts in creating an educational software system, *HyperGami*, that permits students to create and explore polyhedral forms on the computer screen, and to then construct tangible paper versions of those shapes. Our user models have been informed both by systematically observing students using HyperGami and by conducting a study (with 24 undergraduate participants) to illuminate

IUI 97, Orlando Florida USA
© 1997 ACM 0-89791-839-8/96/01 ..$3.50

the difficulties that adults have in understanding and interpreting three-dimensional shapes. In its present state, our user modelling component for HyperGami is only partially complete, and has thus not yet been incorporated in the HyperGami system "out in the field"; ultimately, our goal is to employ this component as part of a larger "critiquing system" for visual thinking (building on ideas of Fischer and his colleagues [12]), in which the computer can advise students of heuristics for visual thinking and interpretation as applied to particular mathematical and design tasks.

Before proceeding further, it should be noted that the status of "mental imagery" (and by extension, "visual/spatial cognition") as an experimental construct in cognitive science has long been a matter of intense debate. Some researchers (e.g., Pylyshyn [23]) argue that there is little support for a uniquely "imagistic" element to cognition that cannot be accounted for by more traditional symbolic models; while others (e.g., Kosslyn [18]) present arguments for an integral role for mental imagery in visual processing tasks such as object recognition. This paper will not seek a resolution of this debate (Gardner [15] and Tye [26] present good chapter- and book-length summaries of the arguments, respectively; while Block [3] includes a number of provocative essays both "pro" and "con" on the subject). Instead, we take an essentially pragmatic approach to the question: our goal is to represent users' understanding of three-dimensional shapes in terms that are primarily visual/geometric in nature (e.g., which faces or edges are "imaged" by the user, and in what orientation), and to present the user with visual-thinking heuristics that make use of these terms. We believe that this "visually-oriented" language carries pedagogical value, inasmuch as it respects people's typical introspective experience of understanding solid geometry; nonetheless, we acknowledge that our visually-oriented vocabulary may ultimately come to be viewed as a suggestive shorthand for some more general (and by implication less "visual") representation.

In the remainder of this paper, we present our ideas for implementing user models for tasks in spatial cognition. The following (second) section motivates our work by presenting a brief overview of the HyperGami system. In the third section, we discuss several experimental results relevant to the question of what makes certain HyperGami solids "difficult" or "easy" for users; and we outline the essential points of a pragmatic computational metric for "degree of difficulty" in visualization. The fourth section describes how these computational ideas have been implemented, to date, as working user models to be incorporated in the HyperGami system; and we conclude with a discussion of ongoing and future directions for research.

HYPERGAMI: AN EDUCATIONAL SYSTEM FOR THREE-DIMENSIONAL DESIGN

HyperGami is written (by the first two authors) in MacScheme [S1] and runs on all color Macintosh computers. The basic idea behind the program is that it permits users to design customized polyhedral shapes, and to view those shapes in a three-dimensional rendering on the computer screen; the program then "unfolds" those shapes into a two-dimensional pattern (a "folding net") which may then be decorated, output to a color printer, and folded into a tangible mathematical model. Figure 1 depicts two of the windows on the HyperGami screen: here, an octahedron is shown both as a three-dimensional solid (in the "ThreeD" window) and in folding net form (in the "TwoD" window). The user is in the process of decorating the folding net, employing a variety of techniques: hand-drawn figures, turtle graphics, solid colors, and patterns have been employed in the decoration shown. (Still other types of decorative strategies—including, e.g., color gradients—are available in the program as well.)

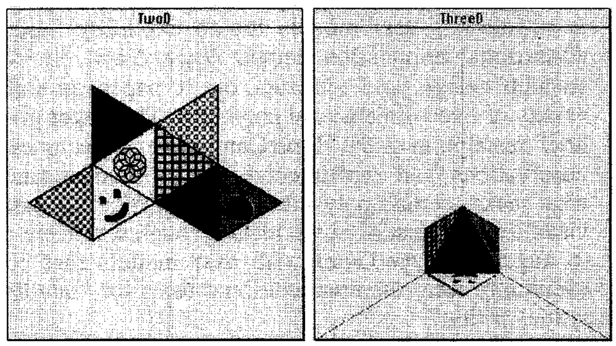

Figure 1. A HyperGami solid (at right, in the ThreeD window) and corresponding folding net (at left, in the TwoD window). The user has decorated the folding net, which will eventually be printed out and folded. The figure also shows another feature of the HyperGami system, in which the user has elected to "paint" the decorations of the folding net over the view of the three-dimensional solid, to provide an indication of what the eventual solid will look like when it is folded.

An especially powerful feature of the HyperGami system is its design as a *programmable application* [8] in which direct manipulation interface techniques are combined with an interactive programming language. HyperGami includes an "enriched" version of the basic MacScheme language—a version augmented with an extensive vocabulary of application-specific procedures and objects. The user can thus employ HyperGami's version of Scheme to quickly create an endless variety of customized polyhedral solids (starting from simpler shapes). Figure 2 shows an illustration: the user has "truncated" the octahedron of Figure 1 at one of its vertices (producing a "jewel-like" shape), and the program has produced a folding net for this new solid. HyperGami in fact contains the ingredients for what might be considered a "functional algebra" of solid objects, allowing users a variety of means for creating new solids from old ones: by combining two solids together at a face, by slicing a solid through a plane, or by stretching (or shrinking) a solid along some axis.

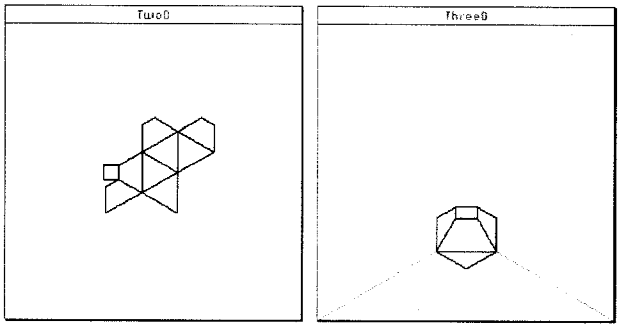

Figure 2. The folding net and three-dimensional solid view of an octahedron that has been truncated at one vertex.

Our interest in developing HyperGami is in creating an exploratory and artistic tool for solid geometry, appropriate both for children and adults; much more detail, both about the system and about its educational uses, can be found in [9]. To date, the system that we have fielded has not been augmented with a user modelling component; but in observing the sorts of difficulties that students have in visualizing three-dimensional solids, we have come to the belief that HyperGami could be usefully extended by modules that anticipate which solids might be problematic for users to visualize. With the addition of such modules, the HyperGami system might eventually be employed to offer advice in how to develop skills of spatial visualization. The remainder of this paper describes our progress in this direction. We begin, in the following section, with a discussion of our observations (and experimental results) in investigating the question of how the sorts of polyhedra constructed in HyperGami are visualized by students.

WHAT MAKES THREE-DIMENSIONAL SOLIDS DIFFICULT (OR EASY) TO VISUALIZE?: SOME EXPERIMENTAL RESULTS

The attractiveness and fascination of polyhedra is often noted by teachers of mathematics [21]; but it is worthwhile to ask what particular factors make certain polyhedra appear "easy to visualize" for students, while others are viewed as more difficult. As an example, we might contrast the two shapes shown in Figure 3: a decagonal prism and a dodecahedron. Both shapes have twelve faces; and the mathematical description of the dodecahedron is, by some measures, simpler than that of the prism (the dodecahedron, for instance, has only one type of face and only one dihedral angle between faces, while the prism has two types of faces and two values for the dihedral angle). Nonetheless, most observers would report that the prism "feels", intuitively, like a simpler shape. Clearly, then, the most obvious metrics—number of faces, types of faces, and so forth—are inadequate to describe the reasons that certain shapes are introspectively deemed "easy" or "difficult."

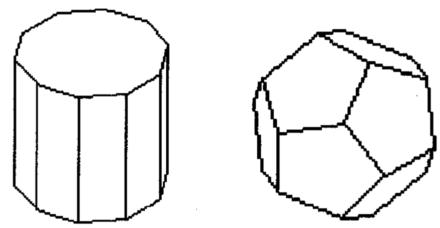

Figure 3. The decagonal prism (left) and the dodecahedron.

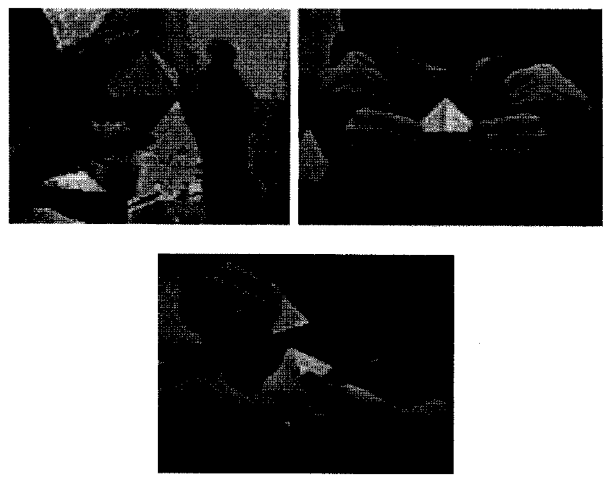

Figure 4. Three elementary-school students holding an octahedron as they discuss their view of the shape. In each case, the child held the shape in a "diamond" orientation (and all three of these children verbally described the shape as a "diamond" as well).

One important clue toward answering this question may be found in the manner in which people place preferred orientations on solids. Consider, for example, Figure 4: this figure shows several stills from videotaped interviews that we conducted with three representative elementary-school-aged HyperGami users during the past year. In these interviews, the children were asked to discuss (among other topics) their ideas about certain solid shapes; in the figure, the three children are shown holding an octahedron (originally shown in Figure 1, and depicted in two alternative orientations in Figure 5). A striking element of the children's discussion of the octahedron is their consistency in its orientation: for each child, the shape is held in a "diamond" orientation (i.e., with a vertex at the top and bottom)—and indeed, all three children described the shape as a "diamond"—even though this is not the orientation that the solid would occupy if it were placed on a table. Figure 5 shows this "stable" configuration of the octahedron.

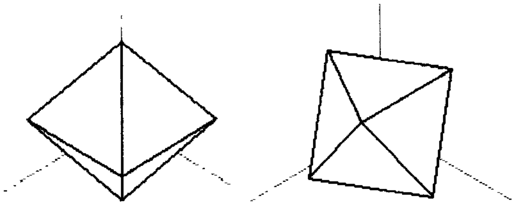

Figure 5. The octahedron, in its "diamond" (left) and "stable" (right) configuration.

The "preferred" orientation suggested by the video stills in Figure 4 was corroborated by a more recent systematic study. In this study, 24 undergraduates (14 males and 10 females) were tested for their interpretations of various solid shapes. Each subject was asked to look at 12 polyhedral shapes: in the experimental procedure, the subject extracted the shape (a paper model) from an opaque box; handled the shape and studied it for 30 seconds; and then replaced the shape in the box. With the shape out of sight, the subject was then asked both to sketch the shape and to write a brief (at most 2-sentence) description for the shape. This protocol was repeated for each of the 12 separate shapes. After viewing all 12 shapes, the subject was also asked to write down a name for each of the shapes (if the subject felt that the shape could in fact be described by a name). (Compare Woodrow's [27] study of children drawing cubes, and Mitchelmore [20] for a study along similar lines conducted with children in Jamaica and the United States.)

Space considerations preclude a detailed description of the results of this experiment ([10] presents much more thorough analysis), but several particular results are worth mentioning here. In sketching the octahedron, 20 of the 24 subjects produced a drawing that could (conservatively) be interpreted as presenting the solid in the same "diamond" orientation indicated in Figure 5; indeed, 10 students included the term "diamond" either in their description or one-word name for the solid (and 9 of those who did not use the word "diamond" indicated an interpretation of the solid as some sort of "double pyramid"). Similar consistent orientation preferences were seen for a number of other shapes: Figure 6 shows several solids in their "preferred" orientations. On the other hand, for the "doubly-capped cube" shown in Figure 7, 16 of 24 students drew the shape in its "vertical" orientation and 8 in its "horizontal" orientation; thus, while there was a clear preference for the former orientation, a significant number of subjects chose an alternative orientation.

The results of this study (and of the earlier interviews) suggest that in forming images of solid shapes, people have a number of orientation heuristics at work. In some cases (as for the octahedron, or the "house" in Figure 6) these heuristics tend to produce an overwhelming preference for a particular orientation (or, what may be the same thing, a rejection of alternative orientations). In other cases (as for the "doubly-capped cube") the heuristics may conflict to produce (among a group of individuals) more than one "typical" orientation.

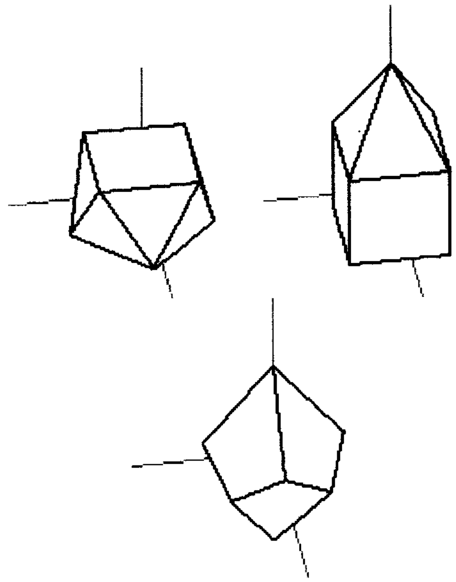

Figure 6. Several solids in their apparently preferred orientations (as viewed from a region in the octant with positive x-, y-, and z- coordinates). The square antiprism (upper left) has its two square faces at "top" and "bottom"; the capped cube (upper right) is configured as a "house"; the trapezohedron (bottom) has its degree-four vertices at "top" and "bottom".

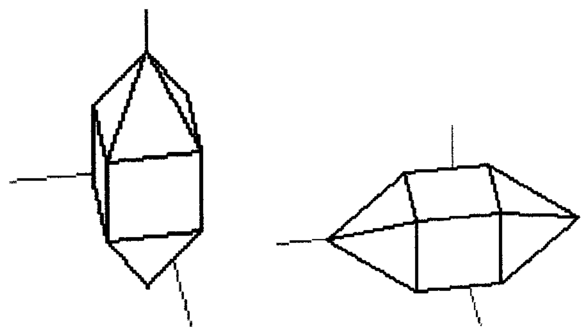

Figure 7. The doubly-capped cube in the two alternative orientations ("vertical" and "horizontal") approximated by subjects' drawings.

Our belief is that an understanding of people's orientation preferences for solids can provide an important insight into why certain polyhedra are seen as "difficult" or "easy." In order to make an estimate of a particular shape's level of "visualization difficulty", we must first produce a plausible preferred orientation for the solid, and then rate this orientation according to the presence or absence of a number of features. Empirically, these "desirable" features appear to include: a preponderance of non-diagonal edges (with some preference to vertical as opposed to horizontal edges); a vertical axis of rotational symmetry; bilateral symmetry; and stability (i.e., the apparent center of gravity of the solid is over either a "flat" bottom face of the solid, or is directly over the bottom-most vertex of the solid). By implication, the "preferred" orientation for a given solid is one that gives rise to a combination (in some sense, the "best" combination) of desirable visual features; likewise, a solid S will be interpreted as "easy" compared to S' if the preferred orientation for S results in a larger combination of desirable visual features than does the preferred orientation for S'.[1]

It would certainly be possible to conjecture innate biological bases for these "desirable" visual features of solids. The human visual system has a higher degree of sensitivity to vertical (as opposed to diagonal) stimuli—this phenomenon is sufficiently well-known to go by the name of the "oblique effect" [2]; thus, a solid orientation (such as that of the "stable" octahedron in Figure 5) in which diagonals predominate may be rejected in favor of an orientation whose edges can be better resolved visually. In a similar vein, human beings themselves (when viewed frontally) exhibit bilateral symmetry, and this may account for a preference for solid orientations that likewise exhibit this feature. (See, for instance, the discussion in Braitenberg [4], pp. 43-47 and 129-130.) In any event, while these would clearly be profitable directions to pursue, our own immediate interests are in using the empirically-determined "desirable visual features" (as inferred from the favored orientations of particular solids) to produce serviceable initial estimates indicating which solids will be viewed as difficult by users. In the following section, we turn to the question of implementing these estimates within the HyperGami system.

INCORPORATING MODELS OF SPATIAL COGNITION IN HYPERGAMI

There are two central tasks involved in producing a computational estimate of the "degree of difficulty" for a given solid S: first, we need to offer a preferred orientation for S (i.e., we need to find some orientation that produces a combination of desirable features); second, we need to evaluate that preferred orientation for S, taking into account the number of desirable features achieved. We now examine, in turn, these two tasks and how they are approached in the current implementation of HyperGami.

Finding a Preferred Orientation of a Solid

We produce a preferred orientation for a solid S in HyperGami by a generate-and-test method in which the program first produces a collection of candidate orientations for a solid, and an "orientation evaluation" module then compares these orientations for desirable visual features. To produce an orientation, the algorithm first assumes that the solid is being viewed from a position along the positive y-axis, and further assumes (by convention) that the center of mass of the solid's vertices will be at the origin; thus, the "top" of the solid is above the xy-plane (i.e., in the half-plane determined by $z >= 0$), and the "left" of the solid (from the viewer's standpoint) is in the half plane determined by $x > 0$. The generation portion of our generate-and-test algorithm now tries two types of rules by which to orient a given solid: by rendering a particular set of parallel edges vertical; or by finding a line between two "important" locations on the solid (where the "important" locations are vertices, edge midpoints, and face midpoints), checking whether that line passes through the center of mass of the solid, and (if so) making that line coincident with the z-axis. Once these rules are applied, "secondary" rotations may be attempted (for instance) to orient a solid so that additional edges will appear horizontal to the viewer (i.e., these edges will be parallel to the x-axis). Often, these orientation rules will in fact produce identical solids; so the algorithm removes duplicate orientations and compares distinct candidate orientations. While these rules do not produce an exhaustive list of orientations, they do appear to produce a plausible set from which to choose a "best" orientation. Figure 8 depicts the four candidate orientations produced for a cube using these rules: the first (at upper left) is chosen by first rendering a set of edges parallel to the z-axis, then rotating the solid so as to produce a set of edges parallel to the x-axis. The second orientation first finds a line between two (opposite) vertices of the cube and makes this line parallel to the z-axis; the system then tries to rotate the cube so that a maximal number of edges are parallel to the x-axis; failing this, the algorithm rotates the cube so that it is bilaterally symmetric when viewed along the y-axis. The third orientation is similar to the second, except with an additional rotation by 180 degrees about the y-axis. Finally, the fourth candidate orientation of the cube has taken a line between two other "interesting points"—two midpoints of the cube edges—and has made that line coincident with the z-axis; subsequently, the cube is rotated so that a maximal number of edges are parallel to the x-axis.

---

[1] Note that these features apply to the solid *in a given orientation* as seen by a hypothetical viewer. For instance, the two orientations of the doubly-capped cube in Figure 7 differ in their visual features: only the (left) "vertical" orientation includes a four-fold axis of rotational symmetry about the vertical axis, while the (right) "horizontal" orientation appears more "stable" in that it rests on a face as opposed to a vertex.

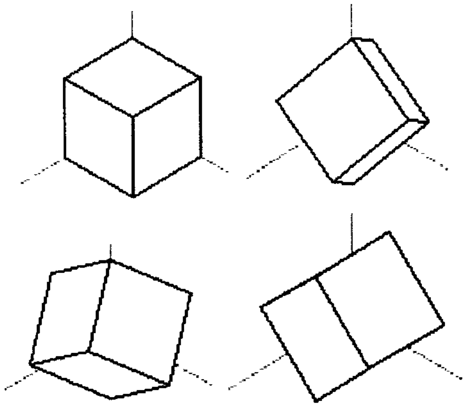

Figure 8. Four candidate orientations for a cube.

The "test" portion of the generate-and-test method for orienting a solid now compares the candidate orientations (such as those shown in Figure 8) for their desirability. For instance, among the orientations for the cube shown in Figure 8, the first is judged to be preferable to the latter three primarily on the basis of the presence of a large number of vertical and horizontal edges; the stability of the figure (the fact that it rests on a flat face) is also noted as a desirable feature of the this first candidate orientation. The algorithm thus deems this orientation the most desirable overall, and this is the orientation of the cube that the program produces.

Figure 9 shows two other examples of "preferred" orientations suggested by the program (along with one other "less desirable" candidate orientation). It should be noted that for these shapes, the program's preferred orientation is at least roughly consistent with the most commonly offered orientation indicated by the experimental subjects described in the previous section: for instance, the program (like the subjects) prefers the octahedron in a "diamond-like" orientation, and the capped cube in the "house" orientation. (It likewise prefers the square antiprism in the "square-on-top-and-bottom" orientation; and it prefers the somewhat more-popular "vertical" orientation of the bicapped cube shown in Figure 7, mildly favoring a pattern in which the "tip" of a pyramid-like portion of a solid is in a vertical line with the center of the solid.) Beyond such coarse-grained qualitative statements, the program's orientations are admittedly hard to compare exactly with those of the subjects, inasmuch as the subjects' responses were produced in the form of a drawing; in general, the human subjects seemed to prefer drawing the solids in a somewhat "oblique" representation whereas the program prefers a more "head-on" view of a solid. Figure 10 offers an illustration of this distinction: the left side of the figure shows a representative view of a triangular prism as drawn by one of the experimental subjects, while the second shows the view of the prism when given its preferred orientation by the program. For both the program and human subject, then, the prism is oriented with a rectangular face toward the bottom; but the program's preferred view is in principle directly "into" one of the triangular faces of the solid, which would therefore be the only face visible.

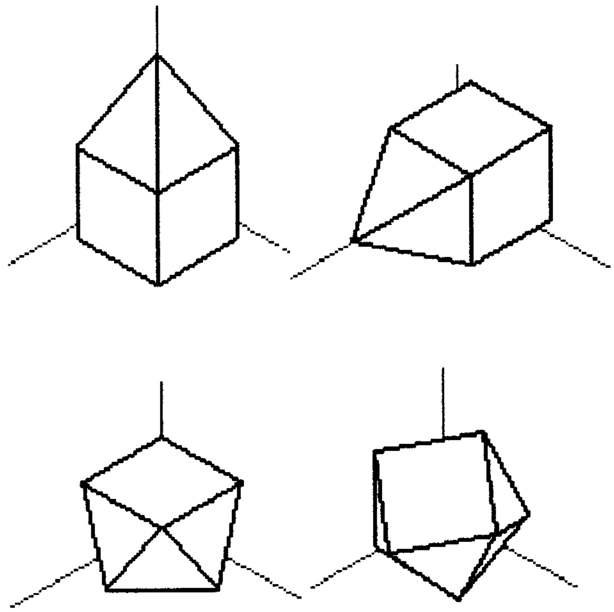

Figure 9. Two shapes with their preferred orientations (at left), along with a sample "less-preferred" orientation (at right). Top row: a capped cube ("house"). Bottom row: a square antiprism.

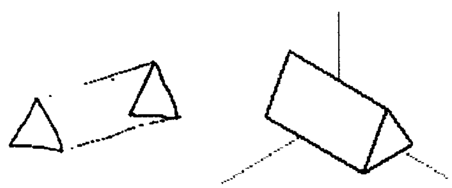

Figure 10. A subject's drawn prism (left) and the program's preferred orientation of the prism (the viewer is assumed to be sitting on the positive y-axis, at the right of the picture, looking toward the origin). The program is thus "looking" at the prism straight along the y-axis and its "view"—unlike that drawn by the student—is simply of a triangle.

Estimating Difficulty of Alternative Solids

The second step in producing an approximate metric of "difficulty of visualization" is to compare distinct solids, in their preferred orientations, for desirable visual features. Once the algorithm described in the previous subsection has produced preferred orientations for two solids S and S', a "solid-comparison" algorithm produces a judgment of relative difficulty between these two solids. This comparison reflects a variety of heuristics similar to those used in judging between orientations of one particular solid: it assumes that an "easier" solid will include fewer diagonal edges, a stable (flat) bottom, and bilateral symmetry. Unlike the orientation-comparison algorithm, the solid-comparison algorithm does not assume that a high degree of rotational symmetry about the z-axis is preferred, nor does it assume that a greater number of faces pointed toward the viewer is necessarily desirable. (Retaining these heuristics would tend to favor, as "easier", solids with similar overall structural properties, but a greater number of faces and edges: for instance, a decagonal prism in its best orientation would be interpreted as "easier" than a hexagonal prism.)

Using this method for estimating "difficulty of visualization", the program produces a partial ordering of three-dimensional solids according to their visualizability. Comparing the best orientations of various solids, we obtain a rank—from easiest to hardest—as follows (see figures 6, 7, and 9):

Cube > House > Octahedron > Bicapped Cube > Square Antiprism > Trapezohedron

ONGOING WORK AND FUTURE DIRECTIONS

While much progress has been made in developing pragmatic user models of tasks that employ spatial cognition in HyperGami, an even greater amount remains unexplored. First, while these algorithms do in fact produce plausible estimates of "the easiest orientation of a solid" or "whether a solid S is more readily visualizable than some other solid S'", they run at a severely slow pace (estimating the best orientation of even a relatively simple solid may take well upward of fifteen minutes), rendering them currently impractical for rapid interactive use in HyperGami. That is, we would not wish a HyperGami user to informally cobble together some new solid and immediately have to pause for twenty minutes while the program produces an estimate of the difficulty of the newly-generated solid. The immediate tasks, then, are to explore ways of generally speeding up the user modelling process—by having the algorithms run as "background tasks" within the application, by optimizing the algorithms themselves ("tuning" them for speed), and perhaps by sacrificing some accuracy or caution in the current algorithms (to obtain even rougher estimations of best orientation and "easiest" solid).

User models of the kind that we have described here may ultimately be of interest in interface design for applications besides HyperGami. Consider (e.g.) what such models might imply for the design of an "intelligent" CAD system. Ideally, such a program could estimate which views of a solid-under-construction are likely to be thought of as difficult or easy, and could accordingly suggest "easy" viewing angles to help the user better understand and visualize a new solid; or it might suggest "hard" viewing angles in the anticipation that these angles would be unexpected to the user. Going a bit further, such a system could make some judgment about whether the user is working on a difficult or simple solid, and could tailor its interface accordingly (e.g., by offering more advice about especially difficult solids, or a wider variety of alternative views).

There are other, more basic—and perhaps more interesting—questions that we can now begin to ask by building upon the user models that we have implemented. We might wish to know, for instance, to what degree the heuristics that we have identified and modelled computationally are culturally dependent (cf. Mitchelmore [20]) : are people from (say) other geographic regions or cultural backgrounds dependent upon the same heuristics, and—if so—do the degrees of dependence differ? (For instance, we might wish to see whether the preference for bilateral symmetry in solids is indeed rooted in human biology as opposed to cultural experience.) We might wish to know whether differential strengthening or appearance of these heuristics appear to account for the development of spatial cognition in children (for instance, in comparing orientations of a given solid, does a preference for stability appear to predate a preference for higher degree of rotational symmetry in children?). And finally—thinking again of the pedagogical purpose of the HyperGami system—we would like to refine our current notions of user modelling to produce truly individuated (and potentially diagnostic) portraits of spatial cognition: that is, we may eventually be able to determine whether an individual has greater or lesser difficulty in considering alternative orientations for solids, and whether practice in this activity appears to impact other measures of spatial cognition. Even better, we may be able to suggest still other techniques for visualizing three-dimensional solids by looking for fresh "spatial heuristics" with which to interpret solids; that is, people might be taught to "look for embedded solids" (such as the cube whose vertices are at the center of each face of the octahedron), or to "look for unexpected symmetries" (such as the three-fold rotational symmetry about each vertex of the cube). In this fashion, we may be able to eventually produce theoretically motivated "mathematical visualization heuristics" analogous to the mathematical problem-solving heuristics advocated by such writers as Polya [22] and Schoenfeld [24].

ACKNOWLEDGMENTS

Our thanks to Mina Johnson and Braden Craig for their assistance in working with experimental subjects. Special thanks to Gerhard Fischer, Clayton Lewis, and members of the University of Colorado's Center for Lifelong Learning and Design for wonderful conversations. This research was supported in part by the National Science foundation under a Young Investigator Award IRI-9258684. Ann Nishioka is supported by a scholarship sponsored by the National Physical Science Consortium.

REFERENCES

[1] Anderson, J. R.; Boyle, C. F.; Corbett, A.; and Lewis, M. Cognitive Modelling and Intelligent Tutoring. Artificial Intelligence, v. 42 (1990), pp. 7-49.

[2] Appelle, S. Perception and discrimination as a function of stimulus orientation: The "oblique effect" in man and animals. *Psychological Bulletin*, 78:4 (1972), pp. 266-278.

[3] Block, N., ed. *Imagery*. Cambridge, MA: MIT Press, 1981.

[4] Braitenberg, V. *Vehicles*. Cambridge, MA: MIT Press, 1984.

[5] Burton, R. R. Diagnosing bugs in a simple procedural skill. In D. Sleeman and J. S. Brown, eds. *Intelligent Tutoring Systems*. London: Academic Press (1982), pp. 157-183.

[6] Burton, R. R. and Brown, J. S. An investigation of computer coaching for informal learning activities. In D. Sleeman and J. S. Brown, eds. *Intelligent Tutoring Systems*. London: Academic Press (1982), pp. 79-98.

[7] Carbonell, J. R. AI in CAI: An artificial intelligence approach to computer-aided instruction. *IEEE Transactions on Man-Machine Systems*. (1970) MMS-11(4): 190-202.

[8] Eisenberg, M. Programmable Applications: Interpreter Meets Interface. *SIGCHI Bulletin*, 27:2, 1995, pp. 68-83.

[9] Eisenberg, M. and Nishioka, A. Creating Polyhedral Models by Computer. To appear in *Journal of Computers in Mathematics and Science Teaching*, 1997.

[10] Eisenberg, M.; Schreiner, M.; and Nishioka, A. A computational model for orientation of solids in spatial cognition. (*Working title; in preparation.*)

[11] Ferguson, Eugene S. *Engineering and the Mind's Eye*. Cambridge, MA: MIT Press, 1992.

[12] Fischer, G., Lemke, A.C., Mastaglio, T. and Morch, A. The Role of Critiquing in Cooperative Problem Solving. *ACM Transactions on Information Systems*, Vol. 9, No. 2, 1991, pp. 123-151.

[13] Fomenko, A. *Visual Geometry and Topology*. Springer-Verlag, Berlin 1994.

[14] Gardner, H. *Frames of Mind*. New York: Basic Books, 1983.

[15] Gardner, H. *The Mind's New Science*. New York: Basic Books, 1985.

[16] Hadamard, J. *The Psychology of Invention in the Mathematical Field*. New York: Dover, 1949.

[17] Johnson, W. L. and Soloway, E. M. PROUST: an automatic debugger for Pascal programs. *Byte*, 10:4 (1985), pp. 179-190.

[18] Kosslyn, S. *Image and Brain*. Cambridge, MA: MIT Press, 1995.

[19] Miller, A.*Visual Imagery in Scientific Thought*. Cambridge, MA: MIT Press, 1984.

[20] Mitchelmore, M. Three-Dimensional Geometric Drawing in Three Cultures. *Educational Studies in Mathematics*. v. 11 (1980), pp. 205-216.

[21] Pedersen, J. Why study polyhedra? In M. Senechal and G. Fleck (Eds.), *Shaping Space: A Polyhedral Approach*, pp. 133-147. Boston: Birkhäuser, 1988.

[22] Polya, G. *How to Solve It*. Garden City, NY: Doubleday, 1957.

[23] Pylyshyn, Z. *Computation and Cognition*. Cambridge, MA: MIT Press, 1984.

[24] Schoenfeld, A. (1985) *Mathematical Problem Solving*. New York: Academic Press.

[25] Siemankowski, F. and McKnight, F. Spatial Cognition: Success Prognosticator in College Science Courses. *Journal of College Science Teaching*, 1:1 (1971), pp. 56-59.

[26] Tye, M. *The Imagery Debate*. Cambridge, MA: MIT Press, 1991.

[27] Woodrow, D. Children Drawing Cubes. *Mathematics Teaching*, v. 136 (1991). p. 30-33.

Software.
[S1] LightShip Software. *MacScheme*. Palo Alto, CA.