ELSEVIER

Information Sciences 140 (2002) 85–103

INFORMATION SCIENCES
AN INTERNATIONAL JOURNAL

www.elsevier.com/locate/ins

TICLE: using multimedia multimodal guidance to enhance learning ☆

Lori L. Scarlatos

Department of Computer and Information Science, Brooklyn College, CUNY, 2900 Bedford Avenue, Brooklyn, NY 11210, USA

Received 17 November 1999; received in revised form 23 December 2000; accepted 18 May 2001

Abstract

Tangible Interfaces for Collaborative Learning Environments (TICLE) explores new ways that multimedia can enhance education without becoming the focus of the educational experience. A TICLE system "watches" as children work together on puzzles and other educational tasks in a physical environment. The system then responds as a "guide on the side", providing the sort of encouragement and prodding that a teacher would. This builds on children's innate love of puzzles, their collaborative tendencies, and their inclination to learn by doing. Although such a system cannot replace a qualified teacher, it can help teachers to motivate and reach more students simultaneously.

This paper outlines the strategies used to create a TICLE system, and describes one implementation: a multimedia multimodal system that responds to children playing with a Tangram puzzle. This system has been installed in the Goudreau Museum of Mathematics in Art and Science, located on Long Island in New York. © 2002 Elsevier Science Inc. All rights reserved.

Keywords: Educational applications; Multimedia human–computer interaction; Intelligent multimodal interaction; Intelligent tutoring; Tangible interface; Ubiquitous computing; K-12 math and science education; Guide on the side

☆ This paper expands on an extended abstract presented at CHI99 [25] (the annual conference sponsored by the ACM's Special Interest Group on computer–human interaction). This work is being supported by grants from the National Science Foundation and the Professional Staff Congress at CUNY.

E-mail address: lori@sci.brooklyn.cuny.edu (L.L. Scarlatos).

PII: S0020-0255(01)00182-7

1. Introduction

Grounding our children with a strong foundation in math and science is essential in the 21st century. Although recent initiatives are helping to strengthen math and science education in the United States, some children, particularly women and minorities, still exhibit an aversion to these subjects. Computer games and web-based activities are not always the best learning tools for these children. For one thing, traditional computer interfaces make it difficult to work with others because only one person can control the mouse at a time. Students who are unsure of their abilities almost always relinquish control to someone who "knows what to do". Another problem is that some very valuable learning activities do not translate well from the physical realm to the virtual. Educators have long recognized that physical problem-solving activities – such as building with blocks, sorting objects, and putting puzzle pieces together – are capable of teaching children a variety of quantitative and spatial concepts [12]. Yet it is not always obvious how to build a tower of blocks on a computer screen. If children could come to see math and science as an extension of these enjoyable physical activities, and begin to experience the thrill of discovery, perhaps that experience would make them more receptive to math and science education [15].

This is the mission of the Goudreau Museum of Mathematics in Art and Science in New Hyde Park, NY, where mathematical puzzles – spread out on tables for people to solve either individually or in groups – play an integral role in the learning process. At this museum they understand that engaging children in learning activities helps them to retain the lessons learned and transfer that knowledge to other situations. Physical activities are natural and fun, and therefore the most effective. By allowing collaboration with peers, the learning is further enhanced.

Although children must make their own discoveries, it is equally important to have a qualified "guide on the side" available to keep students focused, help them over stumbling blocks, and ensure that the right lessons are being learned. This suggests that teachers should be working closely with small groups of students on physical learning activities that they find meaningful. Yet the resources needed for this are scarce. Class sizes in our public schools are increasing; and even when special programs are provided, there are never enough qualified instructors to go around. Thus it is difficult for even the best teachers to effectively guide all of their students through learning activities. This proves to be a key problem at the Goudreau Museum, where one instructor working with 35 visitors on more than 20 puzzles is not always able to provide help when it is needed. When students find that they are unable to progress with a particular puzzle and help is not immediately available, they quickly lose interest and an opportunity to learn has been lost.

Tangible Interfaces for Collaborative Learning Environments (TICLE) provides a promising solution to this problem. With tangible interfaces, everyday objects become input devices. Although a flurry of research activity has recently produced self-sensing objects, items that identify themselves, and even toys that respond to the children who play with them, TICLE is unique in several ways. First, TICLE is able to simultaneously track multiple objects, such as puzzle pieces. Second, TICLE prescribes a general method for uniquely representing the state of a collection of tangible interface objects. Third, TICLE is an educational aid that allows people to focus on the task at hand without having to worry about how to give instructions to a machine. Students may turn to the computer for help and further information, or they may ignore the computer completely and still have an enriching educational experience.

The strength of this approach is that children are not forced to work alone or focus on a computer interface. They are free to work either with their friends or alone; help is available if they want it, but is not forced upon them. Instead of molding the educational experience to conform to the technology, we mold the technology to conform to a desired educational experience.

2. Background

To see what has already been achieved, it is instructive to look at three research areas. First, in the area of educational technology, we see how others have used computers to teach critical thinking skills. Second, we look at the state-of-the-art in tangible interfaces, which enables computers to react to manipulations of physical objects. Finally, because we need to keep track of the relative positions of multiple pieces, we need to look at past work on topological relations.

2.1. Educational technology

Numerous researchers are currently studying ways that technology can enhance learning. Although most computer-aided instruction is designed to be used by one person at a time [28], several research groups recognize the importance of cooperative learning to student achievement. At Georgia Tech, the Learning by Design (LBD) project is getting students to learn science by collaboratively engaging in design activities and reflecting appropriately on their experiences [16]. The Electronic-Games for Education in Math and Science (E-GEMS) group at the University of British Columbia is researching and developing strategies and materials to integrate game-like computer activities with other forms of classroom learning. In one study, they demonstrated that collaborative teamwork motivates girls to try harder and helps them to achieve more than they could individually [13]. Recognizing the limitations of conventional computer interfaces, this group has also experimented with using two mice simultaneously so that no one student retains exclusive control [18].

Yet even when students work in groups, it is helpful to have a knowledgeable instructor standing by. By asking appropriate questions, teachers can help students to think about the problem in appropriate ways and thereby bolster their problem solving skills. When teachers are not available to provide this help, computer programs may be developed to temporarily fill in the gaps. Numerous researchers are studying various approaches to intelligent tutoring, such as [3,11,27]. In all of these cases, the goal is to provide the sort of help that a good teacher would. Many are quite successful, and are able to show an increase in learning. However, these systems are also limited in that they focus on individual learners using computer applications with traditional computer interfaces.

The Epistemology and Learning group at the MIT Media Laboratory has transcended these traditional interfaces, and is consequently conducting some of the most exciting and innovative work in educational technology. With their development of "crickets", devices that make use of programmable integrated circuits and infrared sensors, the group has developed a variety of construction kits that enable children to build things that move, react to stimuli, and communicate with one another [22,23]. The group also sponsors a computer clubhouse where children can work on their own projects using this technology.

All of these groups are making important contributions. Yet different children tend to learn differently [4]. For the children who are not being reached by these efforts, an alternative approach to learning and a novel application of technology to teaching are needed.

2.2. Tangible interfaces

Tangible interfaces constitute a relatively new research area that has grown out of the study of ubiquitous computing [14]. Advances in this area are constrained by the capabilities of sensors and signal processing techniques. It is therefore instructive to examine tangible interfaces in terms of the technologies that they employ.

A sensor may be defined as a device that receives a signal or stimulus and responds with an electrical signal [8]. Different types of sensors are designed to detect different types of stimuli. For example, magnetoresistive sensors can detect the proximity, position, or rotation of the source of a magnetic field. Acoustic sensors detect sound waves within a given spectral range. These sound waves may be either emitted from a source, or emitted by the sensor and reflected back by the source. Alone, they are capable of detecting relative proximity of the sound waves; triangulating results from a pair of acoustic sensors can yield a three-dimensional location. Piezoelectric devices, which generate an electrical charge when subjected to stress, are used to sense bending of materials. They may also be employed by the inertial navigation systems of aircraft in sets of orthogonal accelerometers and gyroscopes [33]. Photoelectric sensors can best detect motion or the presence of an object [1].

Sensors are used to track movement in a variety of ways. For example, sensors can detect head motion and hand gestures in virtual reality interfaces [19]. More recently, researchers in the Physics and Media group at the MIT Media Laboratories have been using the user's body as an electrical conductor for graphical user interfaces [26] and intra-communication among devices on the body [35]. Another way to detect motion of an object is with piezoelectrics, which are commonly found in accelerometers. These are used in a variety of self-sensing devices that respond to motion of the object [17,33]. Yet none of these technologies appear to be appropriate for tracking positions of multiple objects simultaneously.

Several tangible interfaces use tags to uniquely identify objects. A system equipped with the proper sensor can then tell what an object is, and react accordingly. One option is to use infrared, as is done with the crickets [23]. Objects linked to one another (and, ultimately, to a power source) can convey information about how they are inter-linked [10]. Other interfaces use electronic tags that employ radio frequencies [30,34]. In each of these cases, the tagged object must be "shown" or otherwise connected to a sensor at close range. Therefore, position is not considered in these applications.

Computer vision, the traditional choice for tracking objects in an environment, remains an attractive choice due to its low cost (just a camera is needed) and flexibility (any number of objects may be tracked simultaneously). It has been used successfully for motion tracking [7] and, more recently, in the luminous-tangible interface developed by the tangible media group at the MIT Media Laboratory [31]. The main problem with computer vision is that it is subject to occlusion and is sensitive to visual noise in the environment.

2.3. Topological relations

Researchers working on artificial intelligence (AI) and geographic information systems (GIS) have produced a large body of literature on spatial relationships. In the area of GIS, the focus is on either topological or direction spatial relations [20]. Although the 2D topological relations are well defined [5] and are useful for describing relations among spatial entities in GIS, they include relations that cannot occur among solid objects (e.g., intersections) and do not represent *how* two entities meet. Direction spatial relations describe how entities relate to an external frame of reference (e.g., north or south-east) and are therefore not transformation invariant.

AI researchers typically extend the ideas behind spatial relationships for GIS to support reasoning about those relations [6], handle boundary uncertainties

[2] and express changes over time [9]. Although these are worthy ideas, they still do not address the problem of how objects meet.

3. Tangible interfaces for collaborative learning environments

With TICLE, we create learning environments where children can play (and learn) together by using manipulatives in the real world. To achieve this, we have had to address three key pieces of technology. First, the system must track multiple objects in the environment simultaneously using low-cost technologies. Second, the system must uniquely express the state of the physical objects in terms of their topological relations. Third, the system must rely on some strategy for effectively communicating with the learners: guiding without being obtrusive, and helping without giving the answers away. This section describes our approach to these problems.

3.1. Tracking multiple objects simultaneously

Tracking the objects is particularly problematic because of the environment in which it must be done. The technology must work in a noisy room filled with active children, and not be sensitive to any clothing or jewelry that the students might wear. Because scores of children will play with the puzzles daily, mostly unsupervised, the tagged puzzle pieces must also be able to withstand a great deal of rough handling. Finally, the technology must be relatively inexpensive, available, and easy to set up, so that other schools and museums may implement their own tangible interfaces.

Although sensors are getting smaller and cheaper all the time, most continue to be plagued by limitations that make them inappropriate for tracking the positions of multiple object simultaneously. The low cost and flexibility of computer vision makes it an attractive solution for the near term. To overcome the limitations of this approach, I decided to adapt Underkoffler's approach [31] for TICLE.

Reflective patterns adhered to the puzzle pieces uniquely identify those pieces while also indicating their locations and orientations. For many puzzles, a pattern of three colored spots is sufficient: one color marks the center of each piece, while a limited set of other colors forms a code that identifies the piece. Arranged co-linearly, these three spots also form a vector that indicates the orientation of the piece. If two or more pieces are interchangeable in the puzzle, they may have the same reflective patterns. A videocam mounted next to a bank of lights sees these reflective patterns in stark contrast against the rest of the scene, making it easier to eliminate noise from the image and focus on the identifying patterns. The computer vision software can then correlate these patterns with a calibrated set of geometric objects to determine their precise positions. For two-dimensional puzzles, placing the camera beneath the playing surface and viewing the puzzle pieces through a Plexiglas tabletop solves the occlusion problem. This also eliminates the problem of pieces moving out of range of the camera: a proper solution (or partial solution) on the tabletop will always be visible.

3.2. Uniquely expressing topological relations

Knowing where the interface objects is only part of the problem. A puzzle will have at least one solution (or arrangement or state) which the TICLE system will need to respond. The system will also need to be able to recognize partial solutions. In order to respond appropriately to the state (or arrangement) of interface objects, a TICLE system must know what that state is. Each arrangement must have a single representation that is invariant to affine transformations such as translation, scaling and rotation. Likewise, any one representation should not be applicable to more than one physical arrangement of the pieces. The representation must be able to differentiate how a pair of objects relates to one another; after all, there is only one right way to put two puzzle pieces together. Finally, the representation must be able to describe all possible arrangements uniquely, such that each state has exactly one representation.

In describing these arrangements, we are primarily interested in how two puzzle pieces meet. This means that we need to indicate which edges/endpoints/faces are touching. In most cases, we can assume that this will sufficiently describe a puzzle solution. [1] Disjoint arrangements are not considered. Most of the other standard topological relations are of little interest here. Because the pieces are solid, one piece cannot contain another, nor can their boundaries intersect.

For example, suppose that we have a puzzle P made up of a partially ordered set of n puzzle pieces, so that $P = [P_1, P_2, \ldots, P_n]$. Identical puzzle pieces are interchangeable, and may therefore have identical labels (i.e., it is possible that $P_i = P_{i+1}$). Each puzzle piece P_i is a solid polygonal area defined by a sequence of m edges along its boundary, $E_i = [e_{i,1}, e_{i,2}, \ldots, e_{i,m}]$. These edges represent one or more simple polygons [2] with their endpoints listed in clockwise order. Then, for any edge e_j defined by vertices v_j and v_{j+1} the interior (solid) portion of the polygon will be to the right of the vector from v_j to v_{j+1}. This allows us to have holes in the puzzle pieces. For symmetrical puzzle pieces,

[1] If the pieces need to be more precisely aligned, artificial endpoints may be added to the boundary description.

[2] A simple polygon is described by a cyclic sequence of non-intersecting straight line segments (edges) connecting a set of vertices (endpoints).

which may fit together several different ways, edges may have the same label. That way, different arrangements do not need to be defined for different rotations of the piece. For example, in a square piece defined by $E_s = [e_{s,1}, e_{s,2}, e_{s,3}, e_{s,4}]$,

$$e_{s,1} = e_{s,2} = e_{s,3} = e_{s,4}.$$

We can now define the following relationships between P_1 and P_2.

Definition 1. Two puzzle pieces P_1 and P_2 *meet* if at least one edge $e_{i,1} \in E_1$ touches at least one edge $e_{j,2} \in E_2$.

Edge $e_{i,1}$ meets edge $e_{j,2}$ if and only if one of the endpoints of $e_{i,1}$ ($v_{i,1}$ or $v_{i+1,1}$) touches edge $e_{j,2}$, or one of the endpoints of $e_{j,2}$ ($v_{j,2}$ or $v_{j+1,2}$) touches edge $e_{i,1}$. We can show this by considering two possible cases.

Case 1. The edges lie on different lines in space. Then, by definition, these edges will intersect only if the point where these infinite lines intersect lies on both edges. If this intersection point is not any of the edges' endpoints, then the boundaries intersect; they do not meet, and so they are not considered here.

Case 2. The edges lie on the same line in space. If the edges meet at one point, then it must be at the endpoints. If the edges meet at more than one point, then consider the set of points that are touching. Because the edges are finite in length, this set of points must also be finite, with a beginning and an end. And because all of the points lie on a straight line, the bounds of this point set must coincide with an endpoint of one of the edges.

Table 1 illustrates all of the ways that two edges may meet. In this table, edge **a** runs right to left, and edge **b** runs left to right. The columns labeled $\mathbf{a}_1$, $\mathbf{a}_2$, $\mathbf{b}_1$, and $\mathbf{b}_2$ indicate whether that endpoint does (1) or does not (0) touch the other edge. This readily lends itself to a hexadecimal code $(0, \ldots, F)$ indicating how the edges meet. Note that some of the cases are equivalent, depending on which edge one is looking at. Therefore we may impose the following rule: if only the endpoint is touching another edge, then that point must be either (a) the first endpoint on the lower numbered puzzle piece or (b) the second endpoint on the higher numbered puzzle piece. Furthermore, the case in which only the corners are touching (code 6 or 9) need not be represented explicitly, for the relationship will already be represented by two other meet codes (codes 5 and A) as shown in Fig. 1. The column labeled **OK** indicates which relationships satisfy these conditions. This set of 11 possible relationships is complete, in that for any two edges that meet, there is a unique relationship that represents it. Because inaccuracies are likely to arise, it is important to consider some degree of fuzziness in these relations. A broad boundary defi-

Table 1
Representing how edges meet

OK	a_1	a_2	b_1	b_2	Code	Represents
	0	0	0	0	0	disjoint
✓	0	0	0	1	1	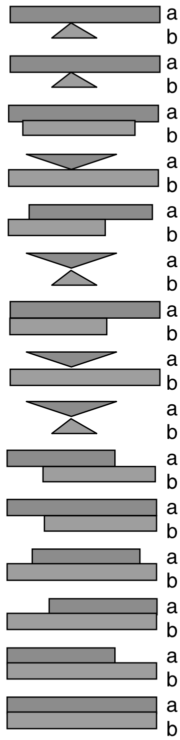
	0	0	1	0	2	
✓	0	0	1	1	3	
	0	1	0	0	4	
✓	0	1	0	1	5	
	0	1	1	0	6	
✓	0	1	1	1	7	
✓	1	0	0	0	8	
	1	0	0	1	9	
✓	1	0	1	0	A	
✓	1	0	1	1	B	
✓	1	1	0	0	C	
✓	1	1	0	1	D	
✓	1	1	1	0	E	
✓	1	1	1	1	F	

nition for the vertices [2] may be implemented by considering a distance ε when examining the proximity of points to edges.

Definition 2. Assuming that puzzle piece P_1 precedes P_2 in the ordering, edge $e_{i,1}$ *touches* edge $e_{j,2}$ if either (a) vertex $v_{i,1}$ is within distance ε of edge $e_{j,2}$, or (b) vertex $v_{j+1,2}$ is within distance ε of edge $e_{i,1}$.

Given the spatial relations defined above, we may represent a puzzle arrangement as a graph with one node per puzzle piece. For each pair of puzzle pieces that meet, a labeled arc indicates how those edges meet. Checking each pair of pieces for a meet relationship can be costly, but is not necessary.

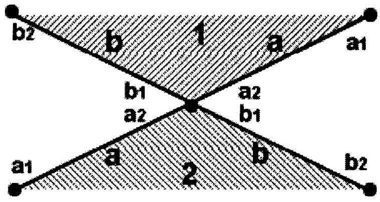

Fig. 1. Two polygons (1 and 2) that meet at the corners are represented by up to four different meet relations.

Two-dimensional pieces may be sorted into two-dimensional bins in linear time, where each bin represents a square region of the space being observed. Size of the pieces determines the optimal size for these squares. Then, the pieces only need to be compared to pieces that occupy either the same bin or adjoining bins. Fig. 2 shows a sample puzzle arrangement and its graph. We can now use this representation to define a solution.

Definition 3. A *solution* is a graph that represents all of the meeting edges in the desired puzzle arrangement. A *partial solution* is a graph that can be made into a solution by simply adding arcs.

It can be shown that the correspondence between the graph and the arrangement of puzzle pieces one to one. Now, each arc in the graph may be uniquely represented by a sub-string of the form:

$P_i.e_i.P_j.e_j.$code,

where P_i and P_j are puzzle pieces $(i \leqslant j), e_i$ is the edge on P_i that touches P_j ($e_i \leqslant e_j$ if $i = j$), e_j is the edge on P_j that touches P_i, and 'code' is a code from Table 1 indicating how the edges meet.

After they are generated, these sub-strings may be sorted in alphabetical order and concatenated together, separated by a ';' character. For example, the following string represents the graph in Fig. 2:

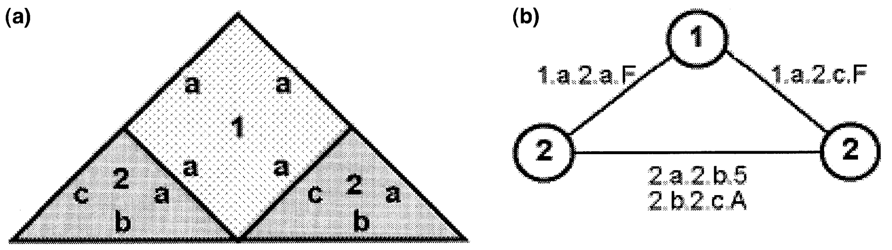

Fig. 2. A puzzle solution (a) may be represented by a graph with labeled arcs (b).

1.a.2.a.F; 1.a.2.c.F; 2.a.2.b.5; 2.b.2.c.A

Because the individual sub-strings are unique, and because an alphabetical sorting is unique, the resulting string is a unique representation of the graph, and therefore of the arrangement itself. For a graph with n arcs, it takes $O(n \lg n)$ steps to generate the sub-strings, sort them, and generate a full string representation. If puzzle pieces are moved one at a time, it takes $O(\lg n)$ time to add an arc to the graph (and a sub-string to the sorted list) and $O(\lg n)$ time to delete an arc.

The current puzzle arrangement is a solution if the string representation of the graph is identical to the string representation of the solution. This may be checked in $O(n)$ time. The current puzzle arrangement is a partial solution if every string representing an arc in the graph is a sub-string of the solution string. If examining the sorted list of solution sub-strings, this comparison takes $O(n \lg n)$ time.

3.3. Communicating with learners

Once the TICLE system knows what the state of the physical puzzle pieces is, it must respond appropriately. Dialog between the children and the computer takes two forms. First, the multimedia feedback lets the children know how they are doing. The trick is to make this feedback helpful without giving the children the answers or overwhelming them with information, and allowing the children to ignore the computer if they want to. This highlights the need for the second part of the dialog: children must be able to ask for help or further information. The students must be able to do this as they continue to work with the puzzle, without leaving it.

Conditions that the system responds to may be placed in the following categories.

1. A solution has been found. The players are congratulated, and the system offers to explain underlying math and/or science principles.
2. A partial solution has been found. The system encourages the players, telling them that they are on the right track.
3. No progress has been made for a relatively long time. The system offers to either give the players a hint or review the rules and goal of the game.

A key factor in this system is determining the appropriate response rate. If the computer reacts to every move every second, it is likely to become annoying. If, however, it waits too long, it may become ineffective.

Multimedia feedback greatly enhances the effectiveness of the system. A color graphics monitor located near the game provides the visual feedback. Audio is used to draw the students' attention from the puzzle to the computer screen. A simple graphical representation of the puzzle's current state provides visual confirmation that the system does indeed "see" what the players are doing. Simple animations with voice-overs provide clear and concise instructions, hints, and explanations. In all of these, it is essential that any audio, text, or graphics be clearly discernible to all players working on the puzzle. Therefore synthetic voices, background music, and extraneous sound effects and graphics are avoided.

Multimodal input makes the system more responsive to the players. A touch screen is ideal because it does not need to be controlled by any one child; all have equal opportunity to touch the screen at any time. Touch areas should be large, or even encompass the entire screen.

The hints need to follow a teacher's method of questioning, as described by Polya [21]. Rather than telling the students what to do, or feeding them a lot of facts, the system poses a question regarding a sub-problem. The children must explicitly ask for the solution to the sub-problem, giving them time to think it over – and try a variety of solutions – beforehand. The hints given must be context-sensitive, helping the children to move beyond the current state of the puzzle.

4. Implementing the Tangram

To test these ideas out, I chose to implement a smart objects interface for an old Chinese puzzle known as the Tangram (see Fig. 3). Five triangles, one square, and one parallelogram, precisely cut from a large square, make up the pieces of this puzzle. Although one may choose to reconstruct literally hundreds of different shapes with the Tangram pieces, the most challenging

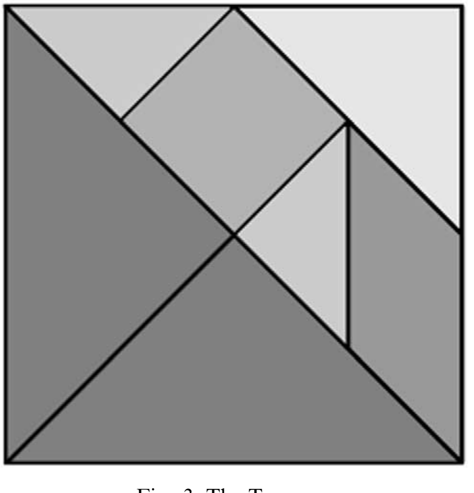

Fig. 3. The Tangram

problem is to reconstruct the square from which the pieces were derived. In solving this initial problem, one may discover underlying geometry principles such as congruence and area.

Others have used the Tangram in many instances for teaching geometry. Tangram problems frequently appear on standardized tests. Pierre M. van Hiele [32], a renowned teacher of mathematics, advocates teaching geometry with a similar mosaic puzzle.

The University of Texas at Austin provides an online Tangram puzzle, along with lesson plans for teachers, as part of its Texas Education Network [29]. The E-GEMS group has been experimenting with different manipulative computer interfaces for the Tangram puzzle [27]. Yet in all of these instances, the traditional computer interface prevents students from collaborating. This reduces the chances that students will discuss the problem, therefore limiting the opportunities for the "thinking out loud" that leads to greater understanding needed for problem solving [21].

For the TICLE-Tangram prototype, I marked the physical Tangram pieces with yellow, orange, magenta, and green reflective spots. Yellow is used solely to mark the centers of the pieces. Identical pieces, such as the two large triangles, have identical patterns of spots. A box-like playing surface was constructed out of opaque Plexiglas, with a frosted Plexiglas tabletop. A bank of fluorescent lights, with a videocam in the midst, illuminates the puzzle pieces from below. A tent structure minimizes the impact of extraneous light from other sources in the room.

I implemented the TICLE-Tangram system on a Macintosh computer using Macromedia Director. This enabled me to rapidly prototype both the interface and the multimedia feedback modules. The TrackThemColors Xtra [24] provides routines for detecting the reflective spots marking the puzzle pieces. An initialization file contains descriptions of the color spots being tracked, the geometry of the puzzle pieces, and the final solution. To account for changes in lighting, the program starts with an option to calibrate the camera before playing. This calibration step simply requires one to click on the reflective spots in an initial video image.

Once playing commences, the system shows a graphical representation of the current state of the puzzle. This is updated every one to two seconds. The players always have the option of either seeing the rules of the game, or getting a hint, by touching a button on the screen (Fig. 4). If no progress is made after a fixed period of time (3 min), the system reminds the players of these options.

When a child requests a hint, the TICLE-Tangram system selects and plays an appropriate multimedia file. Several factors determine which hint will be selected. These factors include the presence (or absence) of certain meet relations, records indicating which hints have already been seen, and whether or not pieces have been missing for awhile. Each hint is posed as a question re-

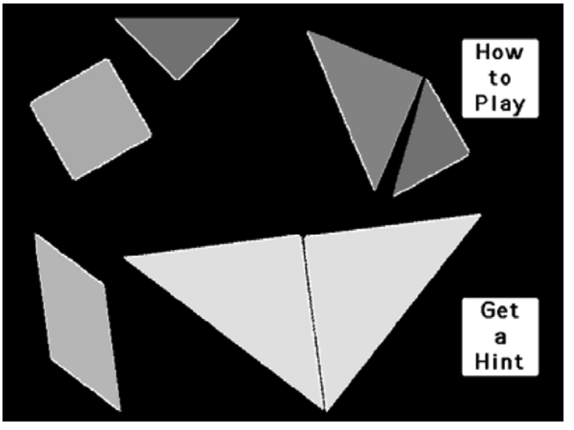

Fig. 4. When the TICLE-Tangram puzzle is in play, the computer display shows the current state of the puzzle, and provides two options to the players.

garding how a subset of the Tangram pieces can be used to make some other shape. The players must then touch the screen to see the solution to the sub-problem. This gives them the chance to think about what is being asked before getting the hint. If progress is detected (i.e., the number of solution sub-strings found increases), audio encouragement is given. When the solution is found, the system congratulates the players. All system actions – including states of the puzzle, and hints played – are recorded in a time-stamped file for later analysis.

This TICLE-Tangram system was installed and tested in the Goudreau museum of mathematics in art and science (Fig. 5).

5. Evaluation

Developing the TICLE-Tangram prototype, we made the following assumptions:

- In a situation where everyone can manipulate the physical pieces at once, the children will be more likely to participate and learn from one another.
- A "guide on the side" – giving hints, encouragement, and reminders about the objectives – can help motivate students and prevent them from giving up too soon.
- Hints that get children to think about the underlying math principles will help them to develop problem-solving skills that will transfer to similar problems.

When we tested our Tangram puzzle system, we hoped that a usability study would provide evidence supporting these assumptions.

Fig. 5. Children working with the TICLE-Tangram prototype at the Goudreau museum of mathematics in art and science.

5.1. Designing the study

We decided to test the Tangram prototype on a group of children who were entering the fourth or fifth grade. There were three test sequences with two groups of three children participating in each test sequence. Children were allowed to select their own groups. In each sequence, one of the groups was asked to work with Tangram pieces being tracked by the TICLE system, while the other group, the control, was asked to work with a standard Tangram puzzle without the benefit of guidance. Both were asked to construct a square from the set of Tangram pieces. Each group was videotaped as they worked. They were given 15 min to complete the task.

When the allotted time was up, the children were shown how to construct the square from the Tangram pieces. They were then asked to put the pieces together in the shape of a house. At this stage the children worked individually, placing small Tangram pieces inside an outline of a house on a piece of paper.

In the final part of the test, students were asked a set of questions regarding both their past experience and their perception of the current experience.

Students who had completed the test, or were waiting to begin, worked on other activities in the museum.

5.2. Observations

The need for a guide on the side was supported by our observations. Initially, all of the groups tried to solve the puzzle without help from an instructor. Groups using the TICLE system worked with the puzzle at least five minutes before asking for help. Eventually all groups requested help, but only those using TICLE actually got it. When the computer gave them hints, we observed the children discussing those hints: part of the meta-cognitive process that is so important to problem solving. We also saw that the children using our computer system got closer to finding a solution (i.e., more of the pieces were placed correctly next to one another), and tended to get less frustrated than those who received no help.

With only one exception, the children chose to work with groups of friends. No group had mixed gender. In general, we found that the girls were more likely to work together and talk about the problem at hand. Boys tended to want to take individual control; in several of the groups, they resorted to "taking turns" with the puzzle pieces. When asked whether they liked working with groups, 44% thought that was the best way to work, while 25% would have preferred to work by themselves.

Even though more than 80% of the subjects recalled seeing the Tangram puzzle before, either in the classroom or on standardized tests, none of the groups succeeded in forming a square. However, the TICLE groups took advantage of the hints we offered and were thereby able to construct more of the solution. We found that the parallelogram caused quite a bit of trouble; many did not notice how rotating the piece could make a difference. Several groups seemed inclined to lay pieces on top of one another, in order to make congruent shapes, but did not seem to know how to use those congruent shapes to solve the puzzle. In general, when the subjects got stuck, they started over, rather than simply trying to rearrange a piece or two. The students working without guidance frequently asked for help (which was not given). All groups grew tired and somewhat frustrated before 15 min were up, although the TICLE groups tended to stick with it longer.

6. Conclusions

TICLE represents a novel approach to using technology in education, in that the technology supplements the educational activity without dominating it. The work described in this paper is unique in several ways. First, TICLE is able to track multiple objects simultaneously. Second, TICLE prescribes a general method for uniquely representing the state of a collection of tangible interface objects. Third, TICLE allows students to focus on the task at hand without having to worry about how to give instructions to a machine.

Although this project focuses on enhancing physical puzzles, the TICLE approach described in this paper may apply to a wide variety of educational experiences. These may include anatomical models that students "dissect" in a biology lab; physics experiments that involve an arrangement of levers and pulleys; molecular models constructed in a chemistry class. It may also be used to "check" the assembly of models, furniture, and even equipment.

Work on this project continues, expanding on the basic ideas presented in this paper. Experiments with alternative sensor technologies are being conducted. At the same time, I am experimenting with a pair of videocams to track objects in three-dimensional space using stereo vision. I am also extending the two-dimensional topological relationships described in this paper to handle three-dimensional puzzles. The ultimate goal of this project is to develop transferable technologies that support rapid prototyping and development of tangible interfaces for collaborative learning environments. This will enable us to build a multimedia authoring environment that allows non-programmers (teachers) to develop their own puzzles or learning activities with tangible interfaces.

Acknowledgements

I wish to thank Bill Wilhelms and Tony Scarlatos for initially suggesting the need for this work, and for introducing me to the Goudreau Museum of Mathematics in Art and Science. Tony also spent a great deal of time helping with the assembly of the learning environment, and provided much helpful feedback during the development of the system. Thanks also go to Beth Deaner, the director of the Goudreau Museum, and her staff for being so helpful and supportive throughout this project. I am grateful to Brian Smith, John Underkoffler, Brygg Ullmer, Josh Smith, and Richard Borovoy at the MIT Media Labs for showing me their own innovative work and brainstorming on possible interfaces for the Goudreau Museum puzzles. I also wish to thank the Brooklyn College students who helped with the testing and have been participating in other aspects of the system's development: Shalva Landy, Julia Breban, Robin Horowitz, and Chanie Sandberg. Finally, I wish to thank the National Science Foundation and the Professional Staff Congress at CUNY for their continuing financial support. This paper is based upon work supported by the National Science Foundation under Grant No. 9984385.

References

[1] Banner Engineering Corp., Handbook of Photoelectric Sensors, 1996.
[2] E. Clementini, P.D. Felice, Approximate topological relations, International Journal of Approximate Reasoning vol. 16 (1997) 173–204.
[3] A. Corbett, H. Trask, Instructional interventions in computer-based tutoring: differential impact on learning time and accuracy, in: Proceedings of CHI 2000, 2000, pp. 97–104.
[4] A. Druin, B. Bederson, A. Boltman, A. Miura, D. Knotts-Callahan, M. Platt, Children as our technology design partners, in: A. Druin (Ed.), The Design of Children's Technology, Morgan Kaufman, San Francisco, CA, 1999, pp. 52–72.
[5] M. Egenhofer, Spatial Query Languages, Ph.D. Dissertation, 1989.
[6] B.A. El-Geresy, A.I. Abdelmoty, Order in space: a general formalism for spatial reasoning, in: Proceedings of 8th IEEE International Conference on Tools with Artificial Intelligence, 1996, pp. 183–191.
[7] J. Foley, A. Van Dam, S. Feiner, J. Hughes, Principles of Computer Graphics, second ed., Addison-Wesley, Reading, MA, 1990.
[8] J. Fraden, Handbook of Modern Sensors, second ed., American Institute of Physics Press, Woodbury, NY, 1997.
[9] D. Gagne, A. Trudel, A topological transition based logic for the qualitative motion of objects, in: Proceedings of 3rd International Workshop on Temporal Representation and Reasoning, 1996, pp. 176–181.
[10] M. Gorbet, M. Orth, H. Ishii, Triangles: tangible interface for manipulation and exploration of digital information topography, in: Proceedings of CHI 98, 1998, pp. 49–56.
[11] N.T. Heffernan, K.R. Koedinger, Intelligent tutoring systems are missing the tutor: building a more strategic dialog-based tutor, in: AAAI Fall Symposium on Building dialogue systems for tutorial applications.
[12] F.P. Hughes, Children, Play, and Development, third ed., Allyn & Bacon, Newton, MA, 1998.
[13] K. Inkpen, K.S. Booth, M. Klawe, R. Upitis, Playing together beats playing apart, especially for girls, in: Proceedings of Computer Support for Collaborative Learning '95 (CSCL), 1995.
[14] H. Ishii, B. Ullmer, Tangible bits: towards seamless interfaces between people, bits and atoms, in: Proceedings of CHI 97, 1997, pp. 234–241.
[15] K. Kahn, Helping children learn hard things: computer programming with familiar objects and activities, in: A. Druin (Ed.), The Design of Children's Technology, Morgan Kaufman, San Francisco, CA, 1999, pp. 224–241.
[16] J.L. Kolodner, D. Crismond, J. Gray, J. Holbrook, S. Puntambekar, Learning by design from theory to practice, in: A. Bruckman, M. Gudial, J. Kolodner, A. Ram (Eds.), Proceedings of International Conference of the Learning Sciences 1998, 1998, pp. 16–22.
[17] G. Levin, P. Yarin, Bringing sketching tools to keychain computers with an acceleration-based interface, CHI 99 Extended Abstracts (1999) 268–269.
[18] J. McGrenere, K. Inkpen, K. Booth, M. Klawe, Experimental design: input device protocols and collaborative learning. Technical Report 96-11, The University of British Columbia, June 1996.
[19] A. Mulder, Human Movement Tracking Technology, Simon Fraser University School of Kinesiology Technical Report 94-1, 1994.
[20] D. Papadia, Y. Theodoridis, Spatial relations, minimum bounding rectangles and spatial data structures, International Journal of Geographical Information Science 11 (2) (1997) 111–138.
[21] G. Polya, How To Solve It, Princeton University Press, Princeton, NJ, 1945.
[22] M. Resnick, F. Martin, R. Sargent, B. Silverman, Programmable bricks: toys to think with, IBM Systems Journal 35 (3&4) (1996) 443–452.
[23] M. Resnick, A. Bruckman, F. Martin, Contructional design: creating new construction kits for kids, in: A. Druin (Ed.), The Design of Children's Technology, Morgan Kaufman, San Francisco, CA, 1999, pp. 150–168.

[24] D. Rozin, TrackThemColors, available online at http://www.itp.nyu.edu/~danny/Xtras.html, 1998.
[25] L.L. Scarlatos, Y. Dushkina, S. Landy, TICLE: a tangible interface for collaborative learning environments, in: CHI 99 Extended Abstracts, 1999, pp. 260–261.
[26] J. Smith, T. White, C. Dodge, D. Allport, J. Paradiso, N. Gershenfeld, Electric field sensing for graphical interfaces, IEEE Computer Graphics and Applications (1998).
[27] E. Shaw, W.L. Johnson, R. Ganeshan, Pedagogical agents on the web, in: Proceedings of Autonomous Agents '99, 1999, pp. 283–290.
[28] B. Shneiderman, Designing the User Interface, third ed., Addison-Wesley, Reading, MA, 1998.
[29] Texas Education Network, Tantalizing Tangrams, available online at http://www.tenet.edu/teks/math/clarifying/clgrd4tangram.html, 1998.
[30] B. Ullmer, H. Ishii, D. Glas, MediaBlocks: physical containers, transports, and controls for online media, in: Proceedings of SIGGRAPH 98, 1998, pp. 379–386.
[31] J. Underkoffler, H. Ishii, Illuminating light: an optical design tool with a luminous–tangible interface, in: Proceedings of CHI 98, 1998, pp. 542–549.
[32] P.M. van Hiele, Developing geometric thinking through activities that begin with play, in: Teaching Children Mathematics, 1999, pp. 310–316.
[33] C. Verplaetse, Intertial proprioceptive devices: self motion-sensing toys and tools, IBM Systems Journal 35 (3&4) (1996) 639–651.
[34] R. Want, K.P. Fishkin, A. Gujar, B.L. Harrison, Bridging physical and virtual worlds with electronic tags, in: Proceedings of CHI 99, 1999, pp. 370–377.
[35] T.G. Zimmerman, Personal area networks: near-field intrabody communication, IBM Systems Journal 35 (3&4) (1996) 609–618.