Computationally-Enhanced Construction Kits for Children: Prototype and Principles

Michael Eisenberg, Ann Eisenberg, Mark Gross*, Khomkrit Kaowthumrong, Nathaniel Lee, and Will Lovett
Dept. of Computer Science, Campus Box 430, U. of Colorado, Boulder CO USA 80309-0430
*Dept. of Architecture, Box 355720, U. of Washington, Seattle WA USA 98195-5720
Tel (first author): +1 303 492 8091, Fax: +1 303 492 2844
Email (first author): duck@cs.colorado.edu

Abstract: Construction kits—toys designed for the building or assembly of physical models—often play an important educational role in children's lives. While such kits have tremendous strengths (e.g., they permit children to build three-dimensional models and to learn through tactile experience), they also have interesting limitations. Traditional construction kits offer little in the way of direct communication with their users—for example, a traditional kit cannot offer a student information or advice about how to proceed in building a model. More generally, traditional constructions—i.e., the models produced—tend to be aesthetically and behaviorally limited. This paper argues that through the use of embedded computation, pieces within a construction kit may communicate with each other, with desktop machines, and with their users; and overall, by integrating construction kits with computation, the educational power and expressiveness of these kits can be greatly increased. As an example of many of the ideas presented here, we describe a prototype of a computationally-enhanced construction kit: a set of speech-enabled alphabet blocks. We conclude by discussing a variety of related research efforts and directions for future work.

1. Introduction

Construction kits—toys designed for the building or assembly of physical models—have a venerable place in the history of education. Perhaps the earliest well-known example is "Locke's blocks", a set of children's alphabet blocks popularized by the seventeenth-century philosopher John Locke [Sutton-Smith, 1986, p. 119]. Architectural blocks flourished in nineteenth-century Germany: one representative catalog published by a Nuremberg merchant circa 1800 listed building sets for castles and walled towns [Harley, 1990]. In America, building blocks were recommended to parents in 1826 as a toy to teach "gentle manners" to young children; and by the late 1860's, Charles Crandall's "tongue-and-groove" interlocking blocks had achieved a popularity that foreshadowed the later success of systems such as Lego [Cross, 1997]. Such kits could apparently be quite extensive: the Museum of the City of New York has in its collection "a log-cabin construction set of nearly one hundred pieces dating from the 1850s." [Harley, 1990, p. 6]

In the past century, the landscape of educational construction kits has blossomed in a variety of directions. Building blocks of various sorts remain a dominant feature of that landscape; but in addition, one might include geometric and polyhedral kits (such as the gorgeous Zometool [P5]), engineering and erector sets, kits of malleable pieces (such as "Toobers and Zots" [P4]), plastic models for assembly (e.g., anatomical models, or scale models of airplanes or cars), chemical modeling sets, and many others. Though collectively these artifacts may be less prominent than, say, computer games and simulations in the modern educational imagination, they clearly command continuing affection from parents, teachers, and children themselves.

This paper is itself a reflection of that sense of affection—it is an extended discussion of the educational role of construction kits, and the means by which that role may soon be enhanced by the integration of computational media. We begin by presenting a brief taxonomy of construction kits; this will serve as a framework for the discussion that follows. In the second section, we describe both educational strengths and weaknesses of construction kits as they are traditionally imagined; and we discuss four large-scale means by which computational technology could enhance construction kit design, preserving (or extending) the strengths of kits while perhaps alleviating their weaknesses. The third section briefly describes a prototype that we have created of a computationally-enhanced construction toy: a set of speech-enabled alphabet blocks. This prototype is intended not so much as an exemplar of ideal design, but rather as an exercise—a springboard for discussion of the technical and engineering issues involved in blending computation into construction kits. Finally, we conclude with a discussion of related work, and a look toward those areas of future research that we believe to be profitable.

A Taxonomy of Construction Kits

While there are conceivably many ways in which to organize the sprawling menagerie of construction kits, there are three major dimensions that seem particularly useful for our purposes (i.e., for the purposes of exploring the role of computers in construction kit design). These dimensions of classification will play a recurring role in the remainder of our discussion.

• *Specificity of constructions*. One broad means of classifying a particular construction kit is according to whether it is intended to produce a single, specific construction or (alternatively) a large range of potential constructions. At the extreme of specificity, we might list assembly kits (e.g., for creating anatomical models of people or animals): such kits are specifically designed to produce one unique construction. At the other extreme, traditional Lego sets [P2] (consisting merely of collections of bricks) are highly non-specific and "free-form": within the constraints of the building elements themselves, the child can construct vast numbers of structures. (It is interesting to note that commercial Lego sets have, over time, moved much more toward the "single-construction" end of this spectrum, as the company markets specialized kits for the creation of particular structures.)

• *Domain specificity*. Some construction kits, such as chemical models, are intended for a relatively narrow (or at least identifiable) academic domain, while others are much more general-purpose—or perhaps playful—in intent. Along with molecular modeling kits, we might list architectural blocks, anatomical models, and polyhedral kits as occupying the "domain-specific" end of this spectrum. At the other end, Toobers and Zots, Tinker toys, and Lego bricks are all far less identifiable in their specific academic content or origins. Note, by the way, that domain-specificity and "constructional specificity" (as described in the previous paragraph) are not quite the same thing: for instance, while molecular modeling sets are highly domain-specific, they are affirmatively not "constructionally specific", in that virtually *any* realizable chemical structure should in principle be constructible with these models.

• *Means of Connection/Construction Materials*. One of the more interesting dimensions along which to classify construction sets is according to the means by which individual pieces are connected or linked. Often simple juxtaposition or stacking is sufficient to connect pieces; in other cases (such as in many traditional erector sets), screws, nuts, and bolts have been employed. Many construction-specific assembly kits have pieces connected (more-or-less permanently) by glue. For still other kits, considerable ingenuity has been shown in devising the means by which pieces are connected together. The Zometool polyhedral connector is an especially impressive example of "linkage engineering"; while it is also possible to have pieces communicate via materials such as string (the "String-a-ma-Thing" kit [P3]) or ballistic objects (as in the disturbingly charming "Cat-a-Pult" kit by A. Ganson, in which successive catapults are linked by a series of far-flung plastic cats [P1]). A consideration of the means by which pieces are connected is inevitably accompanied by considering the materials of which those pieces are made: the metal of erector sets, for instance, suggests the use of screws and bolts as connectors. While the default material for most current kits is plastic (in part, no doubt, for reasons of cost), other materials such as wood, cardboard, and even imitation stone (as in the traditional "Anchor blocks") have been employed.

We will return to these dimensions of classification at several points in this paper; but before proceeding, it should be emphasized that these dimensions are neither intended to be mutually exclusive nor collectively exhaustive categories through which to create a taxonomy of construction kits. Clearly, for instance, a notion such as the "construction-specificity" of a kit is intimately linked to that kit's content domain and construction materials; and in general, we cannot treat any of our suggested dimensions in depth without at least some implicit consideration of the others. Moreover, there are still other (and extremely important) themes according to which construction kits might be classified: the intended audience (e.g., considerations of the age or gender of the children for whom the kit is intended), the cost of materials, the kit's integration with traditional classroom materials and curricula, and so forth.

2. Toward Computational Enhancement of Construction Kits

In this section, we start our exploration of the ways in which computational technology could enhance the value of construction kits. In order to pursue this theme, we begin by a (necessarily brief) discussion of the educational strengths and weaknesses of these kits as traditionally imagined. In other words, why "enhance" construction kits at all? What features need enhancement—and what features need to be improved?

Space considerations preclude a thorough discussion of the educational role of construction kits. Still, we can outline those features of kits that we believe are important in the context of this discussion. On the positive side, there are clear—if perhaps hard to quantify—strengths of construction kits in the larger landscape of educational artifacts. Unlike (for example) material conveyed via text or computer screen, constructed models provide tactile, as well as visual, input for their creator: a polyhedral or molecular model held in the hands affords a kinesthetic perception of three-dimensional structure that is arguably difficult by other means—even including the impressive techniques provided by virtual-reality systems. (Such considerations—relating to the role of tactile perception in education—have likewise long been woven into research surrounding the use of such artifacts as "manipulatives" in mathematical instruction.) In their role as physical artifacts, construction kits can provide the young builder with a natural introduction to the behavior of materials—and to considerations such as friction, stability, and modular design. Beyond these explicitly intellectual or cognitive considerations, there are also intriguing social affordances of construction kits. Constructions and assembled models, for example, are often put on display in homes and classrooms, where they can serve as a spur for curiosity or discussion—and perhaps just as important, where they can serve as a source of pride in craftsmanship for their creator.

On the other hand, there are also limitations (pedagogical and otherwise) to traditionally-conceived construction kits. For example, the educational content of construction kits is in general rather obliquely conveyed. A student who builds a molecular model may well have no idea whether her constructed object corresponds to a real (or possible) compound; a child constructing a polyhedron with a geometric kit may well have little opportunity to "see" any interesting structure or symmetry in the object that he has built. In other words, construction kits provide little in the way of direct communication to their users; and as a result, many opportunities for exploration or reflection could conceivably be lost. And there are still other limitations as well. Domain-specific construction kits, for instance, often can only convey a partial representation of their intended subject matter. Molecular models show a molecule's gross structure, for instance, but not more subtle features such as the distribution of electrical charge (there is nothing in a typical model of water, for example, that suggests the strong dipole moment of the molecule). Similarly, anatomical models cannot adequately show the means by which nerve signals are conducted, or the means by which a bird's wings provide lift. And just as construction kits have social affordances, they likewise have peculiar social weaknesses as well. Most construction kits are (at least in our opinion) aesthetically rather sterile and limited. Assembled constructions tend to be static, unchanging artifacts (with a few notable exceptions, such as robotic Lego constructions). And even for "free-form", non-construction-specific kits, there are few opportunities for personalization: a child cannot actually design or decorate her own pieces.

By considering both strengths and weaknesses of construction kits, then, we may arrive at several fundamental means by which computational media could be beneficially integrated with these kits. By and large, the technological notion underlying these suggestions is that the individual pieces of construction kits may be augmented with small, embedded, individual computers, and that the act of "linking" construction pieces may also correspond to enabling communication between their respective computers. This is of course a non-trivial technological assumption to make, but it is by no means unachievable, as later sections will argue. For the remainder of this section, we present our suggested avenues for bridging computation and physical construction kits.

Communication between linked or neighboring construction elements. One of the major purposes of embedding computation within individual construction pieces is to permit those pieces to communicate amongst themselves. It is this capability that permits a collection of pieces to identify patterns of connection within a complete or partial construction. For instance: by having each "atomic" piece in a molecular model communicate with the pieces to which it is linked by a bond, each piece can, by straightforward algorithmic means, determine the overall molecular structure in which it has been placed. Likewise, a set of architectural blocks could determine whether they have been arranged in an arch; a set of bones in an anatomical model could determine whether they have been correctly placed; and so forth. The ability of an overall construction to identify patterns of connection among its pieces is fundamental to the techniques described in the following two paragraphs.

Communication between construction kits and desktop machines. By endowing construction pieces with the ability to communicate, we likewise enable overall constructions to communicate with desktop machines. To pursue the example mentioned earlier: a chemistry student might thus be able to link a newly-constructed molecular model to her desktop machine and receive information about the construction that she has made (e.g., whether this is a common or even chemically feasible molecule). Likewise, a child who has built a polyhedron might receive specific advice or suggestions from a desktop application about how to visualize symmetry operations in the structure that he has built. In general, the computational power of desktop applications (and, by extension, the World Wide Web) may in this way be brought to bear upon a student's physical constructions.

Communication between constructions and the user. Once construction pieces have been augmented by computation, it might well be feasible (and desirable) to augment them with additional means of identifying or expressing their "current state" to the user. This might be accomplished via means as simple as a flashing light or two, or as complex as an embedded LCD screen. Again, to continue our previous example: a molecule of water might have some means of expressing the electronic distribution of the overall construction by (e.g.) differentially shading the individual atoms of hydrogen and oxygen. Similarly, two nodes of a constructed polyhedron might "light up" simultaneously to indicate that they are interchanged under a particular symmetry operation. Construction pieces might likewise be regarded as potential input devices: turning the previous example around, a student might (e.g.) query whether two polyhedral nodes are equivalent under a symmetry operation by pressing or squeezing them simultaneously. In general, then, the potential development of "user interfaces" for educational construction kits should become a subject of much more active research.

Greater expressive range of pieces and materials. Finally, there are numerous avenues by which we might expand the creative possibilities of construction kits themselves through computational means. Construction pieces might be endowed with simple dynamic capabilities (e.g., one could make a building block that expands and contracts rhythmically along a single dimension). An especially powerful means of expanding the expressive range of pieces would be to permit them to be programmable, in the manner of the programmable Lego brick [Resnick et al., 1996]. It should likewise be possible, with the advent of output devices that "print" in wood or plastic, to permit children to design and decorate their own customized construction pieces, thereby moving construction kits a bit further in the direction of personalized artistic media [Eisenberg, 2002; Gershenfeld, 1999].

In summary, then, we see computational media as a source of techniques by which the traditional educational value of construction kits can be both preserved and strengthened. The advantages of working with physical materials—the acquisition of wisdom "through one's hands"—need not be seen as opposed to the use of computation; while at the same time, both the expressive range and pedagogical (or more broadly, communicative) capabilities of construction kits can be vastly improved. In the following section, we make some of these ideas concrete with a working prototype of a computationally-enhanced construction set.

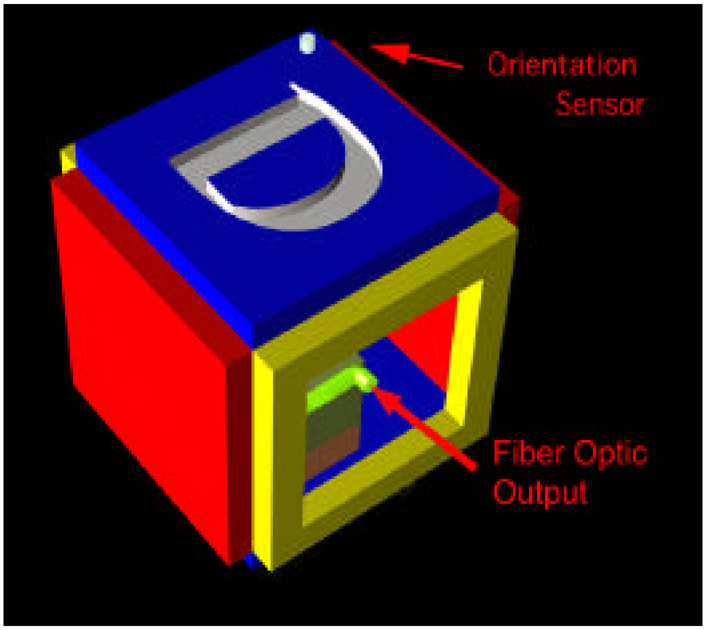

Figure 1. A diagram of the design of a computational alphabet block. One side is open to reveal the fiber optic cable emerging from the edge of the Cricket computer inside.

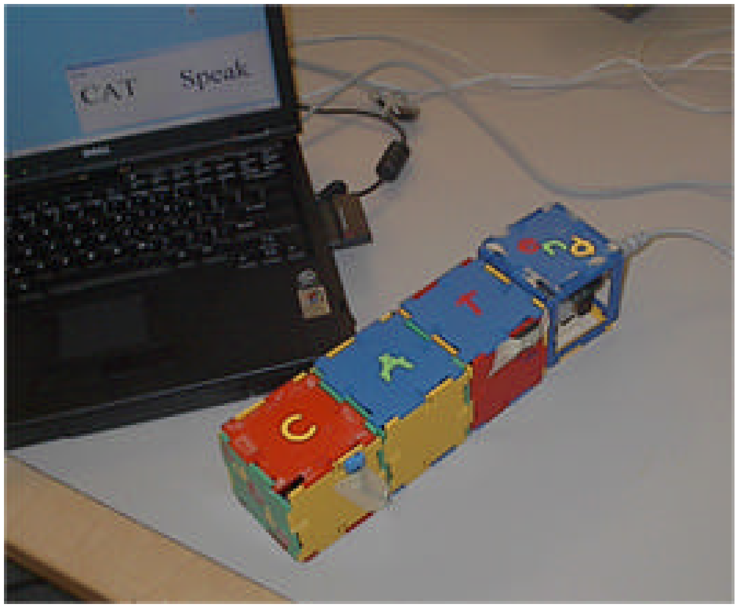

Figure 2. A series of blocks communicates with a desktop machine.

3. Speech-Enabled Alphabet Blocks: a Working Prototype

As an example—really, an existence proof—of some of the ideas discussed thus far, we have constructed a set of "computationally-enhanced alphabet blocks". Essentially, the idea behind these blocks is that they may be arranged in various orders and orientations to construct words; and the pattern of letters so constructed may then be transmitted to a desktop computer, which pronounces the overall word through a speech synthesizer. A diagram of the design of an individual block is shown in Figure 1, and a photograph of a series of blocks communicating with a desktop machine is shown in Figure 2. In the remainder of this section we present an outline of the design and operation of the blocks.

Each block is designed as a simple hollow cube of plastic panels; inside each cube is a small computer (one of the "Cricket" programmable bricks designed by M. Resnick and his colleagues at the MIT Media Lab [Resnick et al., 1996]). When two blocks are juxtaposed so that their letters are in sequence, their respective internal computers communicate via infrared light—each Cricket is equipped with infrared transceivers. For our blocks, the light is transmitted through small holes in the sides of the blocks via fiber optic cables. Thus, for instance, suppose that the letters "C", "A", and "T" have been arranged in sequence. The "C" block communicates a signal indicating its letter to the "A" block that follows it; and the "A" block, in turn, communicates both the preceding "C" and "A" letters to the final "T" block, which communicates the overall pattern of three letters to the desktop machine (which itself pronounces the word "cat"). By arranging arbitrary sequences of physical letter blocks, then, we can communicate (in principle) any desired word to our desktop machine. In the current implementation, touch sensors are used on the exterior of the blocks to determine which of two sides of the block is pointed upward, and which downward.

There are, admittedly, many limitations in this still early prototype—not least of which is that there are only three blocks (and this naturally places a strong limitation on the number of words that we can construct!). In practice, it would naturally be desirable to have a large number of blocks to combine. Another limitation in the current design is that blocks can effectively be only in two orientations—"up" or "down" (i.e., revealing the letter either on the bottom or top of the cube). Thus, in effect, only two of the six sides of each cube can be used as parts of words in our current prototype. This is primarily due to our use of exterior touch sensors to determine block orientation; a better solution would be an internal gravity sensor (presumably along the lines of the sensor used in the "Navigational Blocks" described in [Camarata et al., 2002]), which would permit each block to sense which of the six faces is pointed downward. Finally, in the current design, one of the blocks is designated as a "starter block", sending an initiating signal that is used to indicate the start of a word; in practice, this means that our design requires two "species" of blocks, one for the start of words and one for their continuation. Naturally, an improved design would enable any individual block to be employed in any word position. Still, although the current set is hardly robust enough for anything beyond experimental usage, there is no reason to believe that the basic elements of this design could not be employed and extended to create "block" sets of various types. One variant, for instance, could be to have each block represent a musical note: horizontal sequences of blocks could then represent sequences of successive notes, while vertical stacks of blocks could represent chords. The resulting musical composition could be communicated to a desktop machine (and then played) much as the overall words are communicated (and then spoken) in our alphabetic prototype. It should also be mentioned that our set of blocks could also be extended in the directions suggested by the final two themes of the previous section: currently, blocks communicate with each other and with a desktop machine, but they do not have any means (beyond the exterior decoration) of communicating with the user, nor are they programmable or customizable in any way.

4. Related Work; and a Look to the Future

There are a number of projects that are looking at the potentialities of computationally-enhanced construction sets. Certainly the development of Lego programmable bricks [Resnick et al., 1996] could be described as a major step in this direction; though in practice, the bricks are usually employed as controllers for larger Lego constructions, and not as individual pieces in larger constructions of computational pieces. In the realm of "building blocks", perhaps the most advanced is the ActiveCube project of Kitamura et al. [2001] ActiveCubes are blocks that can be linked into 3D structures which can be communicated to a computer. These cubes are superior to our prototype in that they can be combined in any orientation; and they have advanced embedded input/output facilities as well. In the language of our taxonomy, ActiveCubes are not construction-specific, nor (apparently) are they domain-specific: the examples shown in [Kitamura et al., 2001] simply indicate that constructions may be communicated to a desktop machine, but the blocks do not appear to have associated semantics as (e.g.) letters or notes (although these presumably could be incorporated into particular sets of ActiveCubes). An earlier project along similar lines is reported in Anderson et al.[1999], although these blocks do not have the I/O capabilities of ActiveCubes. Likewise, the gorgeous triangular-tile building set described in Gorbet, Orth, and Ishii [1998] can be viewed as a set of mutually-communicating blocks for creating a variety of three-dimensional structures. One of us (the third author) has collaborated on a project called "Navigational Blocks" in which building blocks may be combined as a sort of input device to construct queries (e.g., individual blocks might represent the "who", "what", and "where" aspects of a particular query). [Camarata et al., 2002] Thus, Navigational Blocks are (at least in their existing instantiation) relatively domain-specific. Yet another domain-specific example is the set of electronic blocks for "tangible programming" designed by P. Wyeth and G. Wyeth [2001]: in this system, computational elements (sensors, logic functions, and actuators) embedded within Lego Duplo Primo blocks [P2] may be combined to form physical representations of simple programs controlling lights, sounds, and wheeled vehicles.

There are a variety of ways of pursuing future research in computationally-enhanced construction kits. We believe that domain-specific kits (e.g., molecular models) represent particularly fertile ground for experimentation, in that pieces may be endowed with domain-specific knowledge. To some extent, our own prototype—as well as the Navigational Blocks, and the Wyeth-and-Wyeth electronic blocks programming kit—are instances of domain-specific kits, although none is perhaps an ideal example: for instance, although our prototype and Navigational Blocks are simple cubes, domain-specific kits are typically designed with geometries suited to their content. (The electronic blocks are only slightly more geometrically differentiated: for the most part, they are simple Lego pieces, but the "actuator" blocks include instances with light bulbs or wheels that indicate their function.) Kits of this type, in which each piece is associated with its own (possibly unique) semantics and geometry, would be plausibly enhanced by embedded computation—e.g., a piece representing a "hydrogen atom" in a chemical modeling set could know something about the chemical properties of the element it represents. Domain-specific kits could moreover include embedded critiquing and coaching facilities to help students reflect upon or understand their constructions. A student's chemistry kit might (for instance) communicate with a desktop machine to suggest pedagogical interventions: perhaps the user should review issues related to bond angles, or perhaps she should explore plausible variants of a particular newly-constructed molecule (e.g., by replacing a hydrogen atom with a methyl group). In contrast, it would be difficult to build such coaching facilities into domain-general building sets, since the kit's embedded software would have less information about the purpose or semantics of a particular construction.

There are still other potential directions for research. One particularly interesting possibility is related to the "material" dimension of our earlier taxonomy: by employing (e.g.) laser cutters or 3D printers to produce customized pieces in wood or plastic, children could create much more personalized constructions. And—going a bit further—by adding end-user programmability to pieces, construction sets could move progressively toward the expressiveness of true artistic media. In short, then, the various projects described here (including our own prototype) should—we believe—be viewed as initial explorations in a much vaster, largely unexplored, and increasingly exciting design space.

References

Anderson D., et al. [1999] Building Virtual Structures With Physical Blocks (Demo Summary). *Proc. of UIST 99*. Asheville, North Carolina.

Camarata, K. et al. [2002] Navigational Blocks. *Proceedings of Intelligent User Interfaces '02*. San Francisco, CA.

Cross, G. [1997] *Kids' Stuff*. Cambridge, MA: Harvard University Press.

Eisenberg, M. [2002] Output Devices, Computation, and the Future of Computational Crafts. *International Journal of Computers in Mathematical Learning*, 7:1, pp. 1-44.

Gershenfeld, N. [1999] *When Things Start to Think*. New York: Henry Holt.

Gorbet, M.G. et al. [1998] Triangles: Tangible Interfaces for Manipulation and Exploration of Digital Information Topography. In *Proceedings of CHI '98*, pp. 49-56.

Harley, B. [1990] *Constructional Toys*. Shire Publications: UK.

Kitamura, Y. et al. [2001] Real-time 3D Interaction with ActiveCube. *In Proceedings of CHI 2001* (Extended Abstracts), Seattle, WA, pp. 355-356.

Resnick, M. et al. [1996] Programmable Bricks: Toys to Think With.. *IBM Systems Journal*, 35:3, pp. 443-452.

Sutton-Smith, B. [1986] *Toys as Culture*. New York: Gardner Press.

Wyeth, P. and Wyeth, G. [2001] Electronic Blocks: Tangible Programming Elements for Preschoolers. In *Proceedings of the Eighth IFIP TC13 Conference on Human-Computer Interaction* (INTERACT 2001).

Products

[P1] *Cat-a-Pult*. HandsOnToys (www.handsontoys.com)
[P2] *Lego*. Lego Company. (www.lego.com)
[P3] *String-a-ma-Thing*. Tobin Toys.
[P4] *Toobers and Zots*. HandsOnToys (www.handsontoys.com)
[P5] *Zometool*. Zometool, Inc. (www.zometool.com)

Acknowledgments

We have benefited enormously from conversations with Gerhard Fischer, Mitchel Resnick, and Carol Strohecker. This work was supported by the National Science Foundation ROLE grant #0106976.