BEHAVIOUR & INFORMATION TECHNOLOGY, 1996, VOL. 15, NO. 1, 57–64

Slow and steady wins the race? Three-year-old children and pointing device use

ERIK F. STROMMEN, GLENDA L. REVELLE, LISA M. MEDOFF, and SETARAH RAZAVI

Children's Television Workshop, One Lincoln Plaza, New York, NY, 10023, USA

Abstract. While adult performance with different pointing devices has received extensive study in the human–computer interaction literature, there is little data on the performance of young children using any input devices at all. In the present study, 64 three-year-old children used a joystick, mouse, or trackball to perform a simple cursor placement task. Two substantive results were obtained. First, trackball users were the slowest, but also the most accurate in their cursor control. Second, characteristics of the children's performance suggest that adult standards for an optimal interface, which stress speed and efficiency, may not be appropriate when children are the intended users.

1. Introduction

The user interface has come to be widely acknowledged as one of the most critical, if not *the* most critical, element in the successful design of computer applications. It is repeatedly stated that 'ease of use' is the key to successful system design (c.f. Myers 1994). But easy for whom? A review of human–computer interaction (HCI) literature reveals that it is dominated by experiments whose subjects are adults, typically executing work-related activities (c.f. Card *et al.*'s (1983) research on document editing, or Wolf's (1992) study of interfaces for executing spreadsheet functions). This research does not focus on the ability of subjects to use a given interface. Rather, the goal of these studies is to determine which interface produces efficient, error-free task performance. Such studies have provided HCI researchers and designers with a rich body of empirical evidence about interfaces for adults using interactive work tools.

Issues confronting child users have not fared nearly so well. It is an intriguing paradox of the field of HCI that interfaces for educational applications are rarely studied in any detail. Despite the much-vaunted educational potential of interactive technologies, few studies in the HCI literature have been devoted to assessments of interfaces for children, particularly children under the age of five. The lack of information about these particular users is all the more striking because it is precisely this population, which is both cognitively and motorically immature (not to mention preliterate), that would seem to be most in need of well-designed, empowering interfaces. The few existing studies do not provide much by way of guidance, but do illustrate the challenges standard interfaces pose to young children.

The serious drawbacks of keyboard interfaces for preschoolers and kindergartners have been described by Grover (1986) and Shade *et al.* (1986), in their naturalistic descriptions of young children's computer use. Comparing the keyboard to both a mouse and joystick, King and Alloway (1992) found that both pointing devices were easier for children to use, although the mouse was best overall. More exotic alternatives have also been considered. A study of speech recognition comparing preschool and adult users revealed that a system that functioned quite well with adults proved to be stunningly inadequate for preschoolers (Strommen and Frome 1993). And while touchscreens appear to have some promise when tested in simple experimental situations (Revelle *et al.* in preparation), serious drawbacks can appear when they are tested in their actual usage contexts (Char 1990). Studies of Nintendo controllers as interfaces also confirm that they pose unique difficulties for young children (Strommen *et al.* 1992, Strommen 1993).

Aside from King and Alloway, however, research on the more common pointing devices and graphical interfaces for computers is generally scarce. Many designers have simply ignored potential problems and adopted the graphical interfaces now common on adult systems for preschool users, employing conventional pointing devices for making onscreen choices. Is the assumption that young children are competent with

0144-929X/96 $12.00

existing input devices accurate? Because these devices have become the *de facto* standards for software meant for young children's use, more information about them is clearly needed. The goal of the present study is to provide more detailed information about young children's ability to use these devices, and to assess the unique developmental issues young children raise as users of educational technology.

The single extant study to specifically compare children's use of all three common pointing devices (joystick, mouse, and trackball) suggests that the assumptions of designers notwithstanding, these devices differ in their ease of use for typical preschoolers (Revelle and Strommen 1990). Their method and results are worth reviewing. In this study, 64 three-year-old children were divided into three groups, each using one of the three devices tested. They were shown how to use the devices to control a cursor in a simple 'electronic colouring book' application, and tested on their ability to place the cursor in specific sections of a drawing and 'colour' it. They were then given 5 days of practice with the software and input device, and tested again on the fifth day.

Two variables were measured: response time, or the average amount of time children required to place the cursor; and error rate, or the percentage of trials where the cursor was not correctly positioned at the time the selection button was pressed. Before discussing the device differences reported, one general finding regarding the children's performance should be noted: the children's response times were extremely lengthy by adult standards, averaging between 9 and 21 seconds *per cursor placement*, depending on the device. This result is similar to those reported in other studies of children's cursor control (c.f. Strommen, 1993), but longer than that reported in psychophysical studies of children's aimed movement (Wallace *et al.* 1978).

In terms of device differences, the results of the study indicated that children's speed with all devices improved over the 5 days of practice. However, notable differences in cursor placement errors were observed. Performance with the trackball was the best on day 1, and did not improve over the 5 days. In contrast, children's performances with the mouse were highly error-prone on day 1, but improved dramatically over the 5 days to end up being as accurate as those of trackball users. Performance with the joystick actually got worse from day 1 to day 5, although the difference in joystick error rates on day 1 and day 5 was not statistically significant. The authors suggested that the trackball's superiority for three-year olds was due to the fact that it provides for the complete mapping of the child's own movements into the rolling of the trackball. While the mouse provides this same mapping, other features of mouse use (learning how to hold the mouse still while the button is pressed, in particular) required practice, and took several days to acquire. The joystick was judged to be the hardest to use because it does not map the child's own movements onto cursor control; only direction of movement is directly mapped onto the cursor. The extra cognitive demands of having to predict and correct for the other aspects of movement that were not under the children's control (distance and speed) were suggested as the reason for the very poor performance of the joystick users.

The results presented by Revelle and Strommen (1990) raise intriguing issues concerning input device use by preschoolers. However, the very global nature of the quantitative variables examined in that study provide only limited insight into children's performance. Consider the dependent variable of average time per placement. Given the substantial lengths of time children required to execute the placements, the question of what exactly is happening during the trials is an important one. Are children simply moving the cursor very slowly, or are there other behaviours apparent in their performance that account for the long times they spend seeking and selecting the target icon? In addition, the definition of cursor placement errors was extremely general. An error was scored if the child actually clicked on any part of the screen that was not the target space. No information on why these errors occurred, or possible factors related to their frequency, was provided. The goal of the present study is to replicate and refine the results reported in Revelle and Strommen (1990). Using more detailed measures, a more complete assessment of young children's use of the same three input devices, a mouse, joystick, and trackball, will be presented.

2. Method

2.1. *Subjects*

Sixty-four three-year-old children participated in the present study. The children were drawn in approximately equal numbers from preschools in Charleston, South Carolina and New York City. Equal numbers of boys and girls served as subjects, and previous experience with computers was assessed via parental questionnaires.

2.2. *Software*

The program used for data collection was created specifically for this experiment. It was designed to be used with a joystick, a trackball, and a mouse. The task was straightforward: move a cursor shaped like Cookie Monster to an icon shaped like a cookie, and press the button so he can 'eat' it. Cookie Monster (CM) occupied a 20×20 pixel square, in a screen environment measuring 200×200 pixels. The 200×200 pixel 'field' on which CM moved was green, and the children were told that it was grass. The rest of the screen space was blue. The children were told that the blue indicated water, where CM could not go because he did not know how to swim.

At the start of each trial, CM always appeared on the screen in a position opposite the location of the cookie. The children were required to move CM across the screen in one of eight different directions (left to right, up to down, diagonal upper left to lower right, etc.) to bring him to the cookie. The movement direction varied randomly subject to a single constraint: all eight different possible directions had to be played out before any direction was repeated. The total possible cookies captured (i.e., the total possible trials) was 40.

To cause CM to 'eat' cookies, children had to move him into the hotspot of the cookie icon, and then press the button on the device they were operating. If he was on the cookie when the button was pressed, he waved his arms, and a brief musical reward played. The cookie then vanished, and a small representation of the cookie appeared as part of a stack, on the far left of the screen in a blue border. The cookies captured during the game session stacked up in this area, and at the end of the session an animation of CM eating them one by one was played for the child, with a musical accompaniment. If, during a trial, a child did not capture a cookie after 2 minutes of game play, that round 'timed out' with an error tone, and a new field appeared. If CM was not on the cookie when the child pressed the button, a 'disappointed' or 'error' tone played. In addition, if CM hit the edges of the field during movement, a different error tone played.

All sessions lasted 6 minutes. At the beginning of each session, children were presented with two practice trials, which were not part of the 6 minute time limit and during which no data were collected, to remind them of the task and of how to use the device. Game play then proceeded.

2.3. *Variables*

2.3.1. *Previous computer/video game experience:* Parents of children in the present study reported children's use of computers or video game machines on a questionnaire. For each type of technology, parents estimated how many times their children had used each system and if the most recent use had been within the past month.

The software recorded five quantitative measures of cursor control competence, which are the focus of the present paper.

2.3.2. *Total time per trial:* This was defined as the length of time from the first movement of the cursor to the cookie being clicked on. This measure was recorded for each cookie captured.

2.3.3. *Time to first target contact:* This was defined as the length of time from the first movement of the cursor to the cursor's first contact with the cookie hotspot. This measure was recorded for each cookie captured.

2.3.4. *Total time the cursor was against the screen border:* This was defined as the amount of time the cursor remained pressed against the borders of the 'playing field', rather than moving in the open space. This variable was recorded for each cookie captured. The more time spent against the walls, the poorer the child's ability to control the cursor.

2.3.5. *Total number of hotspot contacts:* This variable was a measure of accuracy of cursor positioning. It was defined as the number of times the cursor contacted the target before finally stopping on it so it could be captured. This variable was recorded for each cookie captured. The ideal score for this variable is 1.0.

2.3.6. *Total number of cookies captured (total trials)* in the 6 minutes of game play. This was a measure of overall cursor control competence.

2.4. *Procedure*

Children played the game, using the same device, at one session per day for 5 days in a row. All sessions were videotaped for later analysis. The children were tested individually, in a room a short distance from their classroom. In each device condition, the protocol was the same. Children were introduced to the task as follows: 'Do you see Cookie Monster on the computer? Well, this mover makes Cookie Monster walk! Watch! If I do this [action appropriate to device], he walks up. If I do this, he walks down. You try it!' Once the children were using the device, they were prompted to move CM left and right as well. Next, the cookie was pointed out to children, and they were told, 'You see that cookie? Well, Cookie Monster wants to touch it, so he can eat it! Can you make Cookie Monster touch the cookie?' When the children had done so, the Researcher then said 'Now, press the button and watch! He ate it!' More than one explanatory trial was seldom necessary. After several practice trials, the children were then told, 'Now we're going to play a game. You have to help Cookie Monster eat as many cookies as he can before the game is over. Are you ready? Go!'

Table 1. Analysis of variance results for the five performance variables measured in the present study.

		F				
Source	df	Trial time	Time to first contact	Time on borders	Hotspot contacts	Total Cookies
Between subjects						
Device (D)	2	3.42*	23.99***	5.20**	5.55**	0.19
Gender (G)	1	10.01**	0.40	0.82	0.94	11.46**
D × G	2	0.49	4.63**	0.51	0.89	0.52
Error	58	(0.90)	(1.06)	(6.35)	(0.82)	(298.76)
Within subjects						
Session (S)	4	13.94***	7.39***	5.53***	0.34	38.50***
D × S	8	2.21*	0.88	1.05	0.39	2.34*
Error	232	(0.15)	(0.15)	(1.09)	(0.07)	(22.98)

Note: No three-way interactions or Gender × Session interactions were significant, and are not shown. *$p < 0.05$; **$p < 0.001$; ***$p < 0.0001$.

3. Results

3.1. *Data reduction and analysis*

The data from all 5 days of the children's performance were used in the analysis. Except for total number of cookies captured, all the data were standardized across trials for each day by averaging the children's scores across the total trials for that day. The time-based variables and the average number of hotspot contacts were transformed used natural logs to reduce variability. All variables were then submitted to the same analyses: 2 (gender of child) ×3 (device type) ×5 (day of testing) analyses of variance (ANOVAs), where day of testing was a repeated measure. The statistical results of these analyses are shown in table 1.

3.2. *Computer experience*

Eighteen per cent of parents reported that their children had used either a video game or a computer, 10% indicating that such use had happened within the past month. The total number of times of use reported ranged from 'many' (a video game machine) to a single experience. Type of system used, amount of times of use, and recency of use were all not related to any variable in the present study and will not be considered further.

3.3. *Total time per trial*

The ANOVA results indicate significant device effects and day of testing effects, but with a significant device × day of testing interaction. The device × day interaction appears to be due to the finding that while the trackball users remained fairly constant in the length of their trials through the course of the study, joystick and mouse users showed consistent decreases in trial time and continue to get faster across all 5 days.

A main effect for gender on trial time was also obtained. This effect is due to the fact that, on average, girls took longer to capture cookies than boys in all conditions (M = 21.86 seconds vs M = 15.1 seconds for boys).

3.4. *Time to first cursor contact with cookie*

Eight children's data were not included in the analysis because on at least one day, they failed to bring the cursor into contact with a cookie at least once. Five of these children were from the mouse condition; three were from the joystick condition; none were trackball users. It is also notable that seven of these eight children were girls. ANOVA results for the remaining 56 children again indicate significant main effects for device used and for day of testing, along with a significant device × gender interaction. The main effect for day of testing is due to the fact that, with all three devices, over the 5 days children took less and less time to make their first cursor contact (an overall M = 13.45 seconds to first contact with the cookie on day 1, vs M = 6.33 second by day 5). The main effect for device is the result of the children using the joystick being consistently faster than those using the other two devices (M = 7.34 seconds to first contact with the joystick, M = 12.47 second for the mouse, and M = 13.94 second for the trackball).

The significant device × gender interaction appears to be due to the fact that boys made their first contact faster than girls on both the mouse (overall M = 9.36 seconds vs M = 12.62 seconds) and the trackball (M = 10.47 seconds vs M = 15.37 seconds), but girls were faster than boys using the joystick (M = 7.45 seconds for boys, but M = 4.37 seconds for girls).

3.5. *Time on screen border*

Results indicated significant effects for device and for day of testing only. The main effect for device is due to the fact that children using the three devices differed significantly from one another, spending an average of M = 3.82 seconds per trial against the edges of the screen when using the joystick, M = 1.07 seconds when using the mouse, and only M = 0.61 seconds when using the trackball. The day of testing main effect reflects the fact that the amount of time children spent against the walls declined over the 5 days studied, from M = 2.65 seconds on day 1 (averaged across all devices), to M = 1.68 seconds by day 5.

3.6. *Hotspot contacts*

There was only one significant effect for hotspot contacts, a main effect for device group. Joystick users made significantly more hotspot contacts before successfully placing CM on the cookie than did children in the other two device groups (M = 2.54 contacts for the joystick, M = 1.39 contacts for the mouse, and M = 1.25 contacts for the trackball).

3.7. *Total cookies captured*

ANOVA results indicated a significant main effect for session along with a significant device × day interaction. The main effect for day of testing reflects the overall increase in number of cookies captured by children using all three devices (M = 10.84 cookies caught on day 1, increasing to M = 20.39 cookies by day 5). The device × day interaction reveals differential improvement in the different device groups, however. While both joystick and mouse users steadily increased their scores over the 5 days of the study, trackball users started out capturing more cookies on the first 2 days, but their performance then flattened out, such that joystick and mouse users surpassed them by day 5. The day 5 averages for each group were M = 22.13 cookies captured by joystick users, M = 21.19 cookies captured by mouse users, and M = 17.55 cookies captured by trackball users.

A significant main effect for gender was also obtained. Boys consistently captured more cookies than girls on each day of the study and in each device group (overall M = 20.17 cookies for boys vs M = 13.81 for girls.

4. Discussion

The current study provides a detailed look at young children's use of trackball, mouse, and joystick as cursor control devices. The findings confirm those obtained by Revelle and Strommen (1990), but also provide a more detailed portrait of the properties of each device and how they impact on children's cursor control. The present results also raise surprising questions involving larger issues in interface evaluation. The preschoolers' performances in the present study call into question not only the traditional ways we judge pointing devices to be successful, but also whether the adult-based model of what constitutes competent interface use is appropriate when young children are the user population.

4.1. *Is faster always better?*

In studies of adult users, swift performance is almost always viewed as an essential criterion for competent cursor control (c.f. Card *et al.* 1983). A working assumption in the field of pointing device evaluation, in fact, is that devices that conform to Fitt's law concerning the speed and accuracy of performance of the human motor system are optimal for adult use (Mackenzie 1992). Speed, in other words, is assumed to equal excellence. And by this measure, the joystick would seem to be a very successful device for young children. By the fifth day of testing, joystick users were the fastest in making first contact of the cursor with the target, and they and mouse users both captured the most cookies.

Yet by any other measure of cursor control, the joystick is something of a nightmare. The joystick produced significantly more overshooting of the target, requiring extensive repositioning of the cursor by children, than either the mouse or trackball—and this performance did not improve over the five days of practice. Similar results were obtained in a different study of three-year-olds using the joystick-like Nintendo controller, suggesting that overshooting the target may reflect specific developmental cognitive issues (c.f. Strommen 1993). The amount of time children spent stuck on the screen borders reveals a similar pattern. Even though there was improvement over the 5 days of use, the relative ranks of the devices did not change: joystick users consistently spent more time stuck on the screen borders than did users of either of the trackball or mouse.

It appears that these problematic performances did not reduce the joystick's overall time or scores for two reasons. First, because the joystick moves the cursor so swiftly, the extra time children took to repair their errors was often not enough to lengthen their trials substantially. Second, as days of use went by, several children discovered that it was possible to simply keep the joystick button depressed while the joystick was in motion, and that this action would result in the capture of the cookie the moment the cursor intersected it. In other words, they compensated for their difficulty in cursor placement by revising the way they used the device. Obviously, on a screen with more than one 'hot' icon, this strategy would not be useful. In fact, it is possible that the finding by Revelle and Strommen (1990) that placement errors committed by joystick users actually increased over 5 days of practice was due to the children's adopting this behaviour, since in that study the joystick was in fact tested on a screen with many active icons.

If the children's performance using the joystick was fast but error-prone, then their behaviour with both the mouse and trackball could be characterized as slow but accurate. Speed does not distinguish the two devices: children using both devices took approximately twice as long as joystick users, on average, to first contact the cookie with the cursor. What does distinguish the two devices, however, are the changes in children's performances over the 5 days of the study. The trackball stands out as the least difficult device for children to use from the very first day. Compared to children using the mouse or joystick, those using the trackball had the least amount of 'overshooting,' or repeated contacts of the cursor with the cookie prior to clicking. They also spent the least amount of time with the cursor stuck on the borders of the screen. What is more, the relative superiority of the trackball on these measures persisted across all 5 days of the study. Children started out using the trackball more competently from the very first moment, and the relative rank of the devices in terms of cursor control never changed.

The only variable on which the trackball does not appear to be the best device for preschoolers is the total number of cookies captured. But this variable, it turns out, is actually confounded with speed due to the time-limited nature of the task. Speed and number of cookies captured are substantially correlated for each day of use for all devices. In fact, if the time to first contact with the cookie is used as a covariate in the analysis of cookies captured, the main effects for day of testing and for gender remain but the effects for device type and the device $\times$ day interaction both cease to be significant.

What is it about the trackball that makes it the easiest for young children to use successfully? Observations of their manipulation of the device suggest that the answer lies in the physical actions children use to move the ball itself. In a study of children's use of the Nintendo controller, Strommen (1993) found that children's overshooting was greatly reduced, and all aspects of their cursor control improved, if cursor movement was not continuous, but 'stepped', or broken up into short, discrete movements. The steps appeared to be easier for children because they removed the need to execute some of the cognitive schemata necessary to control cursor movement (monitoring cursor position, executing stopping behaviour, etc.), easing the burden on children's working memory.

The trackball appears to be easy for children to use because their use of the device actually produces the same effect as if the cursor moved in discrete steps. Children roll the trackball itself by drawing their fingers across it in a repetitive series of approximately equal-length strokes. At the end of each stroke, the cursor is stationary and its position can be evaluated. The next stroke can then be modified as needed to produce steady movement toward the onscreen target. While no comparative data was collected for this study, informal observations suggest that the speeds with which children roll the ball, and the time they spend evaluating cursor position and computing the direction and size of the next stroke, appear to be longer than equivalent activity when controlling the mouse. In addition, the lack of improvement in trackball users' speed of use after the second day of use suggests that the trackball may impose a ceiling on children's motor performance, simply by the nature of its design. The efficiency lost during trackball use, however, appears to be more than offset by improved accuracy in cursor control.

4.2. *Are we playing the same game?*

The previous discussion focused on the differences in children's performances with the different devices tested in the present study, in order to identify the device children appeared to be able to use most competently. However, regardless of device group, the performances of the children differed from those typically obtained from adults. And while an assessment of the unique features of children's performances was not the goal of the present study, they are worth noting because they raise significant issues concerning the assumptions adult researchers and designers make about young children and computer use.

An obvious quantitative indicator that children's behaviour at the interface differs from that of adults is time: the children's time to first contact of the cursor with the cookie was magnitudes greater than that typically reported in studies of adults using similar devices to perform similar tasks. Why do young children take longer to execute simple pointing tasks? First, there is experimental evidence that in fact, children are systematically slower than adults in executing simple pointing tasks (Wallace *et al.* 1978). Whether this is due to children's immaturity of cognitive processing or the development of motor control is not clear. Whatever its cause, though, children's slower pointing behaviour cannot account for the times observed in the present data: The speeds of children's performances in Wallace *et al.* (1978) were still much faster than those observed in the present study.

The substantially slower times obtained here appear to be due to the fact that the children's conceptions of pointing behaviour and the interface differ systematically from those of adults. The adult attitude can be described as fundamentally pragmatic. For adults, pointing and clicking is simply a means to an end. This orientation has led to the emphasis on efficiency common in studies of input devices with adults, as noted above,. In the present study, the end could be defined as capturing as many cookies as possible. The children, however, do not bring such a goal-oriented orientation to the task. Their engagement of the system was playful, as if the activity was not a task but an opportunity for exploration and self-directed activity. This attitude was reflected in two aspects of their performance that had a direct effect on the length of their trials: manipulating the cursor in idiosyncratic ways, and verbal interaction with the Researcher.

A common behaviour of children in all conditions was to deliberately move the cursor in ways not associated with putting it on the cookie. Children often made Cookie Monster 'dance', or moved the cursor in ways that appeared to be experimental, as if the child was observing the manner in which the device controlled its movement. Prompts intended to bring the child back on task were clearly unwelcome, and only begrudgingly obeyed. In addition to this behaviour, the children would also consistently attempt to strike up conversations with the Researchers while executing the task. Demurrals by the adults had little effect on this performance, and often simply turned the children's attempted dialogues into monologues. The effect of talking on children's cursor movement was to slow it down substantially, apparently because the division of the children's attention between talking and attending to the cursor on the screen interfered with their performance on the task.

It is certainly possible that the task used in the present study, where a cursor shaped like a character is moved on the screen, could have prompted the behaviours noted above. However, it is not the case that the children always 'animated' Cookie Monster as if he was a doll. The effect seems more pervasive than that. In a comparison of two Nintendo interfaces, Strommen (1993) reported that the interface that was notably more difficult for 3-year olds to use was nonetheless more popular than its easier rival, apparently because the children enjoyed the challenge of the interface itself. If the present results can be viewed as more of the same type of data, it suggests that children view the interface in a qualitatively different manner than adults, and that the traditional view that efficiency (i.e., speed of selection) is indicative of a quality interface may need to be abandoned or modified if young children are the intended user population.

4.3. *Conclusion*

The results in the present study suggest that the trackball, which breaks cursor movement into a series of discrete segments because of its physical demands on the user, giving rise to the lowest error rates when positioning a cursor, is the easiest device for young children to use. There is almost no learning curve associated with the trackball either, another indication of its ease of use. However, children's performances with all the devices differed from that of adults in ways that suggest that the traditional adult view of input device use, which assumes that pointing is a fast, task-directed behaviour, may be inappropriate as a standard for this special population. Future research should pursue two different paths in this regard. First, other devices should replicate the current findings regarding trackball use. Confirmation of these results could have significant implications for educational technology design. Second, young children's use of a variety of other interfaces should be examined to determine if the characteristics of their performances reflect those found in the present data. If young children do indeed bring a developmentally distinct attitude to the interface, refining our understanding of this attitude could provide valuable insights into effective interface design for users of all ages.

References

Card, S. K., Morgan, T. P. and Newell, A. 1983, *The Psychology of Human–Computer Interaction* (Lawrence Erlbaum, Hillsdale, NJ).

Char, C. 1990, Touch, click, or jump? The relative merits of different input devises for use by young children. Paper presented at the Annual Meeting of the American Educational Research Association, April, Boston.

Grover, S. C. 1986, A field study of the use of cognitive-developmental principles in microcomputer design for young children. *Journal of Educational Research*, **79**, 325–332.

King, J. and Alloway, N. 1992, Preschoolers use of microcomputers and input devices. *Journal of Educational Computing Research*, **8**, 451–468.

MacKenzie, S. I. 1992, Fitts' law as a research and design tool in human–computer interactions. *Human–Computer Interaction*, 91–139.

Myers, B. A. 1994, Challenges of HCI design and implementation. *Interactions*, **1**, 73–83.

Revelle, G. L. and Strommen, E. F. 1990, The effects of practice and input device used on young children's computer control. *Journal of Computing in Childhood Education*, **2**, 33–41.

Revelle, G. L., Strommen, E. F. and Offerman, S. C. (in preparation). The effects of cursor and device differences on children's computer control. Manuscript in preparation.

Shade, D. D., Nida, R. E., Lipinski, J. M. and Watson, J. A. 1986, Microcomputers and preschoolers: working together in a classroom setting. *Computers in the Schools*, **3**, 53–61.

Strommen, E. F. (in press). Is it easier to hop or walk? Developmental issues in interface design. *Human–computer Interaction.*

Strommen, E. F. 1993, Preschoolers at the interface: a cognitive model of device difficulty, in M. R. Simonson and K. Abu-Omar (eds), *15th Annual Proceedings of Selected Research and Development Presentations, AECT '93*, (Association for Educational Communications and Technology, Washington, DC), 991–1000.

Strommen, E. F. 1992, Formative studies in the development of a new computer pointing device for young children. *Educational Technology*, **32**, 43–51.

Strommen, E. F. and Frome, F. S. 1993, Talking back to Big Bird: preschool users and a simple speech recognition system. *Educational Technology Research and Development*, **41**, 5–16.

Strommen, E. F., Razavi, S. and Medoff, L. M. 1992, This button makes you go up: three year olds and the Nintendo controller. *Applied Ergonomics*, **23**, 409–413.

Wolf, C. G. 1992, A comparative study of gestural, keyboard, and mouse interfaces. *Behaviour and Information Technology*, **11**, 13–23.