The Usability of Handwriting Recognition for Writing in the Primary Classroom.

Janet C Read, Stuart MacFarlane, Matthew Horton

Child Computer Interaction Group
University of Central Lancashire, Preston, PR1 2HE, UK
Tel: *+44 1727 893285*
Fax: *+44 1772 894913*
Email: *jcread@uclan.ac.uk*
URL: *http://www.chici.org*

This paper describes an empirical study with children that compared the three methods of writing - using pencil and paper, using the QWERTY keyboard at a computer, and using a pen and graphics tablet. The children wrote short stories. Where the graphics tablet was used, the text was recognized and presented to the children as ASCII text. Measures of user satisfaction, quantity of writing produced, and quality of writing produced were taken. In addition, the recognition process was evaluated by comparing what the child wrote with the resulting ASCII text. The results show that for the age group considered, writing at the tablet was as efficient as, and produced comparable writing to, the pencil and paper. The keyboard was less efficient. Key usability problems with the handwriting recognition interface are identified and classified, and we propose some solutions in the form of design guidelines for both recognition-based and pen-based computer writing interfaces.

Keywords: Handwriting recognition, Usability, Text entry, Writing, Pen computers, Children, Education.

1 Introduction

This paper presents results from a classroom study to compare three methods of writing with children aged 7 and 8. This experiment is part of a larger investigation that was motivated by a desire to establish whether young children could use handwriting recognition technology, and to identify key usability problems. The context of the work reported in this paper is that of free writing (composition) in the classroom with children aged 7 – 10.

Children spend between 30% and 60% of their school classroom time doing writing activities (McHale et al., 1992). Most of this involves composition work where the child writes down his or her own words in response to some stimulus. Written work is often revised both during and after the initial draft. These revisions include spelling changes, the insertion of missed words, and sometimes, rearrangements of phrases. It is usual for children to limit themselves to revisions that will not ‘mess up’ their work too much. As the process is difficult, it can sometimes happen that a child begins a correction, only to forget what he was going to write by the time he has finished with the eraser (Swanson et al., 1996).

The word processor was designed for revising and manipulating language and so it would appear to have a useful role in written language improvement. Several studies have shown that the use of a word processor can assist in developing written language skills (Sturm, 1988, Newell et al., 1992). A study by Kurth (1987) suggested that, although the use of a computer did not always improve the literal quality of children’s writing, the high visibility of text on the screen fostered more conversation about the writing, and the spelling was improved. Papert (1980) suggests that a computer with word processing software affords the child the luxury of being able to revise and rework their ideas, and therefore becomes a powerful intellectual product.

Children in school are encouraged to use word processors to produce electronic text. However, access to the software is traditionally via a QWERTY keyboard, which novices find difficult to master. Hermann (1987) suggested that handling the keyboard interferes with the writing process, and if this is the case, then the use of more natural interfaces, such as speech and handwriting recognition, may be desirable. The affordability of graphics tablets and pens, and the development of the tablet PC, make handwritten input to the computer a feasible alternative to the QWERTY keyboard.

We begin with a brief summary of research into the usability of handwriting recognition for text input, and specifically into its usability with children. We then describe an empirical study that compares writing outputs from the three methods of pencil and paper, handwriting recognition and keyboard. In the following section we look at recognition rates for both copied and composed text and use this information, together with observations from previous work to derive a list of usability problems and solutions for handwriting recognition technology.

2 Background

Handwriting recognition is the automated process of turning handwritten work into a computer readable form. There are two different technologies, one for off-line recognition and one for on-line recognition. In off-line recognition, the writing is initially captured using traditional means (on standard paper) and is then digitised by scanning technology or by photographic capture; this results in a bitmap or vector image of the writing. The work reported in this paper concerns on-line recognition. Here, the user’s writing is digitally captured at the point of creation;

this is generally done with a graphics tablet or tablet PC and a special stylus or pen. The user's writing is initially displayed as script and is stored as an 'ink' data type. This data type contains information about position and time, and it is this data that is subsequently 'recognised' by the recognition algorithms. This process results in a stream of ASCII characters that is displayed on the screen and is hereinafter referred to as 'transcribed text'. It is common for this text to include errors, as the recognition algorithms generally have to make sense of incomplete or 'noisy' script. Users are likely to include dashes and flourishes and some of their characters will be badly formed.

Research into the usability of handwriting recognition interfaces for text entry tends to focus on the rates of recognition achieved by the human-computer interface and on the elimination or reduction of recognition errors (Read et al., 2001), (Frankish et al., 1995), (MacKenzie et al., 1999). In a handwriting recognition interface, both the user and the system can initiate errors. There have been some guidelines drawn up for the design of recognition based interfaces; these include minimizing the incidence of errors, maximizing their discovery and making recovery easy (Norman, 1981), (Lewis et al., 1986), (Mankoff et al., 1999). Improving the recognition algorithms, constraining the system, and constraining the user can all reduce the number of recognition errors that occur (Plamondon et al., 2000), (Tappert et al., 1990). Constraining methods include limiting the vocabulary, making the user write in boxes, using a reduced alphabet, or using discrete (non-cursive) writing (Frankish, 1999), (Goldberg et al., 1991), (Mankoff, 2000), (Noyes, 2001).

Text input methods are normally evaluated by focusing on the three usability measures of efficiency, effectiveness and user satisfaction (MacKenzie et al., 2002b). Efficiency is typically measured by counting the number of characters entered per second and effectiveness is measured by making comparisons between two strings, these being the presented text (that which was intended to be written) and the transcribed text (that which was produced) (MacKenzie et al., 2002a). For handwriting recognition a CER (Character Error Rate) is derived which is the ratio of the number of errors to the number of characters in the presented text. There have been some longitudinal studies on text input that consider the effect of training and learning (MacKenzie et al., 1999), (Card et al., 1978).

Children develop their handwriting skills in a developmental way with factors such as size, quantity, proportion and spacing all improving with age (Tan-Lin, 1981). Handwriting competence is measured by the legibility of the resulting work and the attainment of a reasonable speed. People are often judged by their handwriting, and it is not unusual for a child with poor handwriting to develop a poor self-image (Sassoon, 1990).

Handwriting may be 'joined up' (cursive) or it can be 'printed' (discrete). It has traditionally been the case that English schoolchildren learn to write using discrete writing, and then move onto cursive writing at around age eight or nine. There is considerable support in the research community for a change to this pattern, with researchers suggesting that children should use cursive script at a much younger age. Peters (1985) view is that children should learn cursive writing as early as possible in the school curriculum, because:

- It results in children understanding the concept of a word at an earlier age
- It results in better letter formation
- Children don't have to 'adapt' at age eight
- The movements involved in cursive writing produce better spelling.

This last claim is particularly interesting; Peters claims that '*quality of handwriting is highly correlated with spelling attainment*', a relationship supported by others including Bearne (1998), who relates the connection between handwriting and spelling to kinaesthetic memory.

We have published on the use of handwriting recognition for text input by children (Read et al., 2002a), (Read et al., 2003a). This earlier work focused on the general usability of the technology and its usefulness to support the writing process. How children deal with and recover from errors, how the technology compares to other text input methods and the mental models that children have of recognition technology have all been reported (Read et al., 2002c), (Read et al., 2001), (Read et al., 2003b).

3 The Empirical Study

The empirical study reported in this paper was designed to determine the relative usability of handwriting recognition for extended writing, compared with the two more traditional methods of pencil and paper, and keyboarding. It was hypothesized that children would prefer to use the pen and tablet and that it would result in more writing being done than at the keyboard. An earlier study, reported in Read et al. (2001) compared mouse, speech and keyboard input with handwriting and concluded that children found the handwriting recognition interface satisfying, and that it compared favourably with the keyboard in terms of efficiency. This earlier study used very short samples of writing and it was not possible to make any judgements about quality.

3.1 Design of the Experiment

Earlier work by the authors had established that a short training session with the handwriting recognition interface produced much better recognition rates. Training for the pencil and paper was considered unnecessary, but it was felt to be important to check familiarity with the keyboard and so a simple task was prepared which established whether the children could find and use the space bar, the full stop and capitalization. For those who had problems, guidance was given.

3.1.1 Design of the Handwriting Recognition Training Interface

This interface presented text in a single line at the bottom of the screen for the children to copy. An alternative to presenting the text visually would have been to present the text in audio, via headphones or loudspeaker. This was considered to be a bad idea as the children using the technology would be unlikely to be able to spell the words fluently and this would cause errors, and increase the children's anxiety.

It was decided to use three phrases in the training, on the assumption that the first would be badly recognised, the second would be significantly better and a third would act as an insurance against spurious results for the first and second. From the third phrase a 'recognition rate' would be derived that could be used in the analysis of the results. Using three phrases also allowed the identification of repeated poorly formed letters. There were two possible approaches to the design of the text to be copied. One was to use the same phrase three times; the other was to provide similar but different phrases. In earlier work, the latter approach was taken as it reduced the effect of familiarity and boredom, and allowed for a meaningful activity; however, for this experiment it was decided to use one repeated phrase as it was considered that this would help the child see what was happening and help construct a more useful mental model. As the child was only writing three phrases (as opposed to ten in the earlier study) boredom was seen to be less of a problem. Additionally, there was a meaningful activity following this training exercise; this had not been a feature of the earlier work.

The training text to be copied was 'a big dog climbed down the tree'. It comprised 25 characters and 6 spaces. The writing fitted on one line on the computer screen in a large font.

There were a number of ways in which the recognized text could be presented to the children. Errors could have been highlighted by the use of wiggly lines under the text as is common in many word processing applications. It was also possible to extract confidence scores for each recognized character, and these could have been displayed or interpreted to give the child some clues about difficult recognitions. These two methods were not used as it was felt that they may cause too much anxiety for the children and they did not necessarily reflect the real world task.

3.1.2 Design of the Writing Task

The recognition software had a dictionary that could be turned on or off. The dictionary is used during the recognition process to allow for word matching – this can result in better recognition rates, although its usefulness with children, who may use non-standard spellings, is less obvious. We decided that, as the dictionary was likely to be present in most applications, having it turned on was the most appropriate model for this work.

Children were given a writing task in line with normal English primary school activities. Each task had a stimulus and an activity. The children were given 12 minutes to write, and they were warned when they had 2 minutes remaining in case they wanted to edit their work. Spell checking of the final text was not enabled in either the recognition-based application or on the keyboard; this meant that if a word was spelt incorrectly but was accurately recognised, the child would not see an error. Children writing by pencil and paper were also not given spelling support.

It is common for children to make errors when they write. These are sometimes discovered during the writing activity, and sometimes noticed later. In all three writing methods, children were given the opportunity to correct errors.

3.1.3 Subjects

Eighteen children aged between 7 and 8 were recruited from an English Primary (age 5 – 11) School in Lancashire. The class teacher selected the children, choosing a mixed sample of abilities and gender, but each child came to the test as a volunteer and was given the opportunity to leave the research activity both before and during the work. All the children had English as their first language and none had used handwriting recognition-based technology, or the pen and tablet before. There were seven boys and eleven girls; two of the children were left-handed.

3.1.4 Apparatus

Three story stimuli were prepared with assistance from the class teacher. These were identified as S1, S2 and S3. Each child taking part in the experiment was identified by a code (C1 – C18) and the input modalities were identified as P, K, and T (for pencil and paper, keyboard and tablet).

For the handwriting application, an experimental interface (Figure 1) was constructed using Visual Basic 6® and the Calligrapher® recognition engine. This was presented to the children on a standard laptop that had a Wacom® graphics tablet attached to it.

The children used the pen that was supplied with the tablet to construct their writing. Prior to beginning the experiment, children had the equipment explained to them. They then carried out the training activity as described above.

The keyboard interface comprised a regular word processing package with a standard QWERTY keyboard. The font size was preset to 14 and the spell checking and grammar checking was disabled. Prior to doing the experiment children were asked to type a short phrase to ensure that they could use capitals, the space control and punctuation.

Figure 1 - The Writing Interface

The pencil and paper interface comprised a pencil, a piece of lined paper and an eraser. No training was given!

3.1.5 Procedure

The design was within-subjects single factor with three conditions: writing using paper, writing at the keyboard, and writing with the tablet. A 3 x 3 Latin Square was used to determine the order in which children did the three activities, and within this a further 3 x 3 square was used to determine which of the three writing stimuli each child used.

Children attended the study in groups of three, with each child being directed to one of the three input methods on entering the room. This meant that each cluster of three children had the same experience with regard to the story stimulus and the environmental conditions. The room that was used was a storeroom with a large table in the centre of it and the children sat on three sides with the two researchers seated between them. On entering the room, each child was asked to indicate on a smileyometer (Figure 2) (Read et al., 2002b) how good they thought the activity to which they had been directed was going to be. The rationale for this was that this gave a measure of expectation that could indicate whether or not the child was subsequently let down by the activity, or pleasantly surprised.

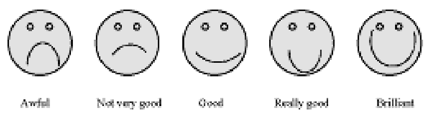

Figure 2 - Smileyometer used to record children's opinions

Following any training, the three children were presented with one of the writing stimuli, and then given ten minutes to write their story. After ten minutes, they were told to stop and given the opportunity to edit their work if they wanted to. They were then asked to rate the input modality for user satisfaction using a new smileyometer to give a rating for 'actual' experience.

For each activity, measurements were taken of

Quality of the writing (Teacher assessed)

Quantity of the writing (Word count)

After the children had done all three input modes, they were asked to assess the relative user experience with the three modalities. Two other tools that have been developed by the authors, a fun sorter and an again-again table (Read et al., 2002b) were used for this. The fun sorter required the children to rank the three writing methods in order of preference, and the again-again table invited them to indicate whether or not they would like to use each method again. This was presented to the children on a single sheet, and was given out a week after the study.

3.2 Results of the Study

The class teacher measured the quality of the writing; the results can be seen in Table 1. All the stories were collected, and made anonymous in such a way that the teacher did not know what method had been used for the writing, nor which child had done the work. Initially, each piece of text was awarded a grade of 1, 2c, 2b, 2a or 3c (as used in the National Literacy Strategy (Dfee, 1988)) and these were then converted to numbers between 1 and 5 where 5 represented a piece of work at level 3c. Nationally children of the age under investigation were expected to be working at a level 2b, i.e. a score of 3.

	Mean	Standard Dev
Tablet	3.0	1.0
Keyboard	2.9	1.1
Paper	3.7	1.2

Table 1- Quality of writing

The quantity of writing (Table 2) was measured using a word count. This was generated automatically using the tool in Microsoft Word®.

	Mean	Standard Dev
Tablet	58.8	25.5
Keyboard	44.0	24.7
Paper	72.6	38.1

Table 2 – Quantity of writing

The results from the two smileyometers were converted into numbers between 1 and 5 where 1 represented the choice ‘awful’ and 5 represented the choice ‘brilliant’. These were averaged across the children and the resulting scores, for both the expected fun and the actual fun, are shown in Table 3.

	Expected	Actual
Tablet	4.2	4.3
Keyboard	4.6	4.7
Paper	4.1	4.3

Table 3 – Children’s preferences

A week after the event, each child completed the short pictorial questionnaire that comprised the fun sorter and the again-again table. The results from these are summarised in Table 4.

Hardest to use	Tablet
Easiest to use	Paper
Most fun to use	Keyboard
Least fun to use	Paper

Table 4 – Results from survey after all three events

4 Discussion of the Results

4.1 Quality of the Writing

The best writing was from the pencil and paper. The stimulus for this writing was a parable about an ant and a grasshopper in which the ant stored food for the winter, while the grasshopper played his cello all summer and subsequently went hungry in the winter. The children were asked to write a story about what happened the next year.

4.2 Quantity of the Writing

The children wrote more using the tablet than they did on the keyboard (N = 17, t = 4.7, P = 0.0001). When writing at the keyboard, children spent a lot of time looking for the right key to press. It was interesting to note that children using the keyboard asked for spellings whilst those using pencil and paper and the tablet did not. Given that these were the same children, it can be hypothesised that the keyboard creates more anxiety or more uncertainty about spellings. The tablet and the pencil and paper both support word-based construction as opposed to letter-based construction.

4.3 Child Preferences

The children marginally preferred the keyboard to the other methods, however, the results were not significantly different. We observed that, when the children were completing the smileyometers, some commented about the quality of their own work (their story) and so it was difficult to be sure about what the child was evaluating, whether it was his or her own work or his or her experience with the interface.

We had expected that the children would prefer the pen and tablet to the keyboard, as it was new to them. It later transpired that only one child had used a word processor for writing before, for all the other children, this too was a novel experience although they were all able to use the keyboard, albeit it slowly..

Children commented after the three activities that one thing that they did not like about the tablet interface was that it didn’t allow them to write their entire story on one screen. When the children wrote on paper and on the keyboard, all their work was visible in the same place, whereas with the handwriting interface, their work appeared both to the right of the writing surface (as ASCII text – and with errors) as well as on the writing surface. It is clear that the children found the tablet hard to use; given that they had only a short time to practise, this was not a surprise.

5 Discussion of the Key Findings

In this section we examine some of the key findings with respect to the fit between the technology, the user, and the task. We introduce some data about recognition rates, and explore the errors that occurred at the tablet interface, and investigate the potential of the three technologies for writing.

5.1 Recognition Rates and Experimental Design

While the children were composing on the tablet, one of the researchers copied their writing (including spelling errors) onto a notepad in order to create a corpus of presented text. There may have been some errors in this text as the researcher had to make some assumptions (for example she had to sometimes guess what character the child was forming), and also may have inadvertently missed a word or misspelled a word. This presented text was then compared with the transcribed text (stored by the software) to derive a character error rate for the text that the child had composed (MacKenzie et al., 2002a). The three phrases of the copied text that had been used for training were similarly analysed. The results are shown in Table 5.

	Mean	Standard Dev
Training phrase 1	50.4	25.0
Training phrase 2	34.8	23.2
Training phrase 3	41.2	20.4
Composed text	27.8	15.7

Table 5 – Error rates

This table (Table 5) demonstrates that error rates fall significantly between the first and second training text. What is interesting from table 5 is that the error rates for the composed text were considerably lower than the error rates for the (copied) training phrases. We suggest that there are two main reasons for this effect; these are discussed in the next two subsections.

5.1.1 The Occurrence of Spaces

When the children are copying text, many spaces are introduced into the generated text due to the time taken by the children to complete words. This is caused by a feature built into the recogniser that inserts spaces between words when there is a 'pen-up' of a specified duration. For adult usage, this is an essential and useful feature, but for children who may be copying words character by character it is a problem. Thus, if a child writes '`climbed`' too slowly, it can be recognised as seven one-letter words rather than as one seven-letter word. This, and the resulting effect on error rates, is demonstrated below

Presented text `climbed`
Generated text `c l i m b e d`

MSD = 6 (represents 6 inserted spaces)
CER = 46.15% (from 6/13 which is the number of errors divided by the number of characters in the transcribed text))

It is therefore likely that copied text will result in higher error rates than composed text.

5.1.2 Focus of Attention.

The term *focus of attention* (FOA) was introduced by MacKenzie et al. (2002b) to describe the attention demands of a text input task. For an expert typist, copying text from a piece of paper is a single FOA task, as the typist needs to look only at the text to be copied and not at the screen or the keyboard. If children are composing on paper with a pencil, they similarly have only one FOA. When children are composing at a keyboard or using the writing tablet they have two FOA – the stylus or the keyboard, and the computer screen. When children are copying text, an extra FOA is created in each case. This overburdens the feedback channels and therefore makes the task more difficult and more error prone.

5.2 Improving the Usability of the HR Interface

Observations made during this and previous experimental work have identified a range of usability problems with the handwriting recognition interface. We have previously classified four types of errors that arise when the child is composing writing at the interface (Read et al., 2001). These are presented in Table 6 together with some new suggestions for solutions to the problems.

In addition to these errors there are some usability issues that relate to the fit between children writing and the use of the interface. These 'task conformance' issues concern behaviours that will be expected from the children, which are not well supported by the handwriting recognition interface. We have identified four such behaviours and we discuss them below.

	Example	**Solution**
Spelling error	Child misspells words	Phonic spellcheckers and phonic dictionary support
Construction error	Child doesn't form the letter or word correctly; 'a' may look like 'd' or 'd' may be constructed in the wrong order	Training for the child Some flexibility in the recognition algorithms
Execution errors	The child doesn't touch the tablet with the pen when he writes	Child to look at the screen

	The child cannot correctly position the pen on the writing space.	Child to look at the screen
	The child continues writing once the pen has turned into a pointer (this happens when the pen moves away from the writing area onto a menu area, it is discussed in (Read et al., 2002d).)	Haptic feedback Child to look at the screen
	The child writes too slowly for the recognition software – introducing spaces.	Training Cursive writing Changing the recognizer 'pen-up' time lag
	The child may be writing outside the writing capture area.	Haptic feedback Child to look at the screen
Software error	The software misrecognises the word or character.	Better algorithms

Table 6 – Errors at the interface and some solutions

5.2.1 Rubbing out / Scribbling out

When children write, they commonly revisit the last word they have written and scribble it out (or rub it out) before substituting another word. Scribbling out has a disastrous effect on recognition, so rubbing out needs to be encouraged and supported! This requires an eraser (often the other end of the stylus) with a different set of behaviours attached to it than those attached to the pen.

5.2.2 Starting at the Top of the Page

Children start to write in the extreme top left hand corner of the page and seem to want to completely fill a line before moving down. The effect of this is that sometimes the pen stroke goes off the writing space. If menus are placed at the top of the screen, they are likely to be opened during the writing activity. It is prudent to put all the controls at the bottom of the screen when children are writing. As writing so close to the edge is not a trait that children should carry into adult life, it is acceptable to encourage them to start writing a little way down the page and to write within right and left margins. This can be implemented by putting a greyed border around the writing area, thus encouraging the children to keep away from the edges.

5.2.3 Fixing Bad Writing

Children are encouraged to produce 'good' handwriting in school. As the children's writing is normally assessed as an end product, it is quite usual for children to go back to their writing and 'correct' it so that it looks better when visually examined. An example of this is given below in Figure 3 where the child saw 'she' after it was written and, realising that the ascender on the 'h' was a little short, went back and added more ascender. The recognition software finds this difficult to deal with, particularly where there is a time lag involved. It is likely that the correction below will be recognised as snel

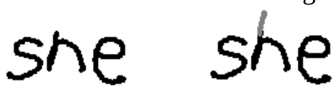

Figure 3 – The effect when children go back and correct their writing

The recognition software could cope with this behaviour but it needs complex algorithms that can be made spatially aware. A delayed stroke solution is commonly applied when a late dot over an i or a late stroke through a t is encountered in handwriting. This could be extended to short ascenders and descenders, to poor hoops and to open joins, all of which children commonly correct on detection.

5.2.4 Adding Missing Things

This is a similar problem; the children go back and insert words, or insert letters into words. This cannot be easily coped with by recognition algorithms. The only way that this can be effectively managed is by identifying these corrections as 'mark up', at which point the recognition can behave in a different way. For this to happen, there needs to be a different pen that has different functions. This could be a physical device or could be a virtual instance of the same pen. With this enabled, the writer could add words and phrases between existing words and phrases with the additions being 'placed' in the correct place rather than occurring at the end of the writing (as determined by time sequence).

5.3 Writing Quality and Digital Ink

As can be seen from table 2 above, the children did not write huge stories. Some children were unable to fully develop their plots in the time allocated. Only a few read over their work when it was done. We had expected that children might edit their work, but they didn't. Research indicates that editing and revision is uncommon at this age (Latham, 2002). What was apparent was that children were able to write quite fluently at both the tablet and the paper.

Writing is typically broken into three overlapping phases, these being

- Planning and collecting

- Initial drafting leading to more final writing
- Revising and editing

(Hayes et al., 1986)

In all phases of this process, the use of tablet technology and digital ink is supportive for children writing. In the first phase, pencil and paper have traditionally been used for note taking, and for storing and recording creative ideas and thoughts. A word processor allows for no annotation and, as we have identified in this study, children become over concerned with spellings. Notes for subsequent writing, sketches, and diagrams can be made almost as easily with the pen technology as with paper.

The second phase does seem to favour the keyboard input, as it will result in reasonably reliable ASCII text. However, children need to break their words into characters, and it is possible that this has a negative effect on the writing that is produced. In addition it is slow to use for children who have not been taught keyboarding, and there is therefore an extra focus of attention. With a tablet PC, like pencil and paper, the child needs only to focus on one thing.

The third phase of this process, revising and editing, is greatly facilitated by the use of computer aided text processing. Traditionally, this has been enabled by word processing software manipulating ASCII text, but it may be more beneficial for children to preserve their handwritten work and manipulate this as digital ink. In table 6 above, the majority of the execution errors can be eliminated if a tablet PC is used instead of a separate graphics tablet. This is a result of there being a reduced focus of attention when the pen and the screen are joined. The errors that are caused by the child having to focus on two locations at once are eliminated.

6 Conclusion

The results reported in this paper have implications for the design and evaluation of handwriting recognition based technologies and digital ink technologies.

The first finding is that once children get over the usability problems of the tablet technology, they can write fluently using a stylus. This bodes well for tablet PCs and for digital pens, as they not only reduce the foci of attention to one, (for better error reduction and better ease of use) but they may also reduce the number of execution errors caused by the pointer behaviour, and the separation of the pen and screen. We are currently conducting a user trial with digital pens and tablet PCs in the writing classroom to investigate these claims.

The second finding is that the digital ink technology supports cursive text construction, which has been shown to improve the spelling of children (as does the pencil and paper). Where the digital text has an advantage over pencil and paper is in its ability to be either recognised and thereafter treated and manipulated as ASCII text, or to be manipulated in its own right as digital ink.

An observation that is of interest to the writing community is that text that is constructed at the computer interface can be logged over time, resulting in rich data for research and development. It is now possible to track the multiple changes that a child (or an adult) may make to a piece of writing.

We have identified some of the problems with using copied text for the recording of recognition rates and error rates. It is likely, where children are involved, that error rates with copied text will always be higher than for composed text, due to the additional focus of attention, and the difficulties that children have in remembering the exact composition of the words to be copied.

Future work will investigate the usefulness of digital ink for writing and the usability of the tablet PC and the digital pen for child users.

References

Bearne, E. (1998) *Making Progress in English*, Routledge, London.

Card, S. K., English, W. K. and Burr, B. J. (1978) Evaluation of mouse, rate-controlled isometric joystick, step keys, and text keys for text selection on a CRT, *Ergonomics,* **21,** 601 - 613.

Dfee (1988), The National Literacy Strategy as a Framework for Teaching,DfEE, London.

Frankish, C. (1999) Pen Input In *Interface Technology: The leading Edge*(Eds, Noyes, J. M. and Cook, M. J.) Research Studies Press, Baldcock, Hants, pp. 59 - 72.

Frankish, C., Hull, R. and Morgan, P. (1995) Recognition Accuracy and User Acceptance of Pen Interfaces *CHI'95 Denver* ACM Press / Addison-Wesley Publishing Co 503 - 510

Goldberg, D. and Goodisman, A. (1991) STYLUS User Interfaces for Manipulating Text *UIST'91 South Carolina* ACM Press 127 - 135

Hayes, J. R. and Flower, L. S. (1986) Writing research and the writer, *Amercan Psychologist,* **41,** 1106 - 1113.

Hermann, A. (1987) Research into writing and computers Viewing the gestalt *Annual meeting of the Modern Language Association San Fransisco, CA* ED 292094

Kurth, R. (1987) Using word processing to enhance revision strategies during student writing activities, *Educational Technology,* **27, (1)** 13 - 19.

Latham, D. (2002) *How Children Learn to Write*, Paul Chapman Publishing, London.

Lewis, C. and Norman, D. A. (1986) Designing for Error In *User Centred Systems Design*(Eds, Norman, D., A and Draper, S. W.) Lawrence Erlbaum, Hillsdale, NJ, pp. 411 - 432.

MacKenzie, I. S. and Chang, L. (1999) A performance comparison of two handwriting recognizers, *Interacting with Computers,* **11, (3)** 283 - 297.

MacKenzie, I. S. and Soukoreff, R. W. (2002a) A Character-Level Error Analysis for Evaluating Text Entry Methods *NordiChi2002 Aarhus, Denmark* ACM Press 241 - 244

MacKenzie, I. S. and Soukoreff, R. W. (2002b) Text Entry for Mobile Computing: Models and Methods, Theory and Practice, *Human-Computer Interaction,* **17,** 147 - 198.

Mankoff, J. (2000) Providing Integrated Toolkit-Level Support for Ambiguity in Recognition-Based Interfaces *CHI'00 The Hague* ACM Press 77 - 78

Mankoff, J. and Abowd, G. (1999) Error Correction Techniques for Handwriting, Speech, and other ambiguous or error prone systems *Interact 99*

McHale, K. and Cermak, S. (1992) Fine motor activities in elementary school: Preliminary findings and provivional implications for children with fine motor problems, *American Journal of Occupational Therapy,* **46, (10)** 898 - 903.

Newell, A. F., Boothe, L., Arnott, J. and Beattie, W. (1992) Increasing literacy levels by the use of linguistic prediction, *Child Language Teaching and Therapy,* **8,** 138 - 187.

Norman, D., A (1981) Categorization of Action Slips, *Psychological Review,* **88,** 1 - 15.

Noyes, J. (2001) Talking and writing-how natural in human-machine interaction?, *International Journal of Human-Computer Studies,* **55,** 503-519.

Papert, S. (1980) *Mindstorms, Children, Computers and Powerful Ideas,* Basic Books, Great Britain.

Peters, M. L. (1985) *Spelling: Caught or Taught: A New Look,* Routledge and Kegan Paul, London.

Plamondon, R. and Srihari, S. N. (2000) On-line and Off-Line Handwriting Recognition: A Comprehensive Survey, *IEEE Transactions on Pattern Analysis and Machine Intelligence,* **22, (1)** 63 - 84.

Read, J. C., MacFarlane, S. J. and Casey, C. (2001) Measuring the Usability of Text Input Methods for Children *HCI2001 Lille, France* Springer Verlag 1 559 - 572

Read, J. C., MacFarlane, S. J. and Casey, C. (2002a) Designing a Handwriting Recognition Based Writing Environment for Children *8th International EARLI SIG Writing Conference Staffordshire, England*

Read, J. C., MacFarlane, S. J. and Casey, C. (2002b) Endurability, Engagement and Expectations: Measuring Children's Fun *Interaction Design and Children Eindhoven* Shaker Publishing 189 - 198

Read, J. C., MacFarlane, S. J. and Casey, C. (2002c) Oops! Silly me! Errors in a Handwriting Recognition-based Text entry Interface for Children *NordiChi 2002 Aarhus, Denmark* ACM Press 35 - 40

Read, J. C., MacFarlane, S. J. and Casey, C. (2002d) Pens Behaving badly - Usability of Pens and Graphics Tablets for Text Entry with Children *ACM UIST'02 Paris, France*

Read, J. C., MacFarlane, S. J. and Casey, C. (2003a) 'Good enough for what? Acceptance of Handwriting Recognition Errors by Child Users *IDC 2003 Preston, UK* ACM Press 155

Read, J. C., MacFarlane, S. J. and Casey, C. (2003b) What's going on? Discovering what Children understand about Handwriting Recognition Interfaces *IDC 2003 Preston, UK* ACM Press 135 - 140

Sassoon, R. (1990) *Handwriting: A New Perspective,* Stanley Thornes, Leckhampton.

Sturm, J. (1988) Using computer software tools to facilitate narrative skills., *The Clinical connection,* **11,** 6 - 9.

Swanson, H. L. and Berninger, V. W. (1996) Individual Differences in Children's Working Memory and Writing Skill, *Journal of Experimental Child Psychology,* **63,** 358 - 385.

Tan-Lin, A. (1981) An investigation into the developmental course of preschool/kindergarten aged children's handwriting behavior, *Dissertation Abstracts International,* **42, (4287A)**.

Tappert, C. C., Suen, C. Y. and Wakahara, T. (1990) The State of the Art in On-Line Handwriting Recognition, *IEEE Transactions on Pattern Analysis and Machine Intelligence,* **12, (8)** 787 - 808.