Based on a paper entered in the 1989 ESU competition. Prepared as an HTML document by John Higgins, June 1997. Click on this link to download a copy of the Hopalong software (DOS only).

HOPALONG: A COMPUTER READING PACER

by John Higgins, University of Stirling
and Ruth Wallace, University of Bristol

Background: the nature of reading skills

It is conventional to distinguish at least three main types of reading: intensive reading, extensive reading, and scanning. This paper is concerned with the second of these, and the term 'reading' in what follows can be taken to mean extensive reading unless it is otherwise qualified. The type of reading most often found in language classrooms is intensive reading or occasionally scanning; extensive reading is not usually trained but seems instead to be left as something to be picked up. It is ironic that activities intended to provide help with extensive reading, particularly the use of a class 'reader', are often subverted into intensive reading by being evaluated. Simply by telling students to read a chapter and to be prepared to answer questions about it, you turn extensive reading into 'study', which is virtually a synonym for intensive reading. Study requires the following:

exact comprehension of sentences
recall of detail and possibly verbatim recall
(a book we have studied is a book we can quote from)
re-reading

and therefore must be carried out at relatively slow speed. When you read extensively, by contrast:

you tolerate some fuzzy understanding
you ignore a great deal of detail
you demand fresh material
("No thanks, I've read that.")

All this makes the reader aim at higher speeds. Some evidence of how reading can be turned into 'study' is on show in every British Council overseas library. The first twenty pages or so of much of the classic fiction will be heavily annotated with marginal glossings, but from the fiftieth page onwards there will be no annotations. Rather than read the whole book, students have studied as much as they could before their time or interest ran out.

Good and bad readers

It is possible to learn quite a lot of a foreign language without acquiring an extensive reading skill or learning to read at speed. However, tertiary level studies demand a good deal of reading and the foreign learner who is only capable of laborious study-mode reading will be hampered. The other argument for trying to develop reading skill is that the skill itself becomes a means of securing the input which enables other language skills to develop. Beginners, of course, are incapable of reading; their unfamiliarity with the language patterns makes them unable to predict and therefore limits them to study. Once learners are past the beginner stage, we can observe whether they turn into good or bad readers.

The frustration cycle

Christine Nuttall (1982:167) provides an interesting description of what she calls 'weak readers' who are trapped in a 'vicious circle' of not enjoying reading, reading slowly, not reading much, understanding relatively little, and therefore not enjoying reading. Good readers, on the other hand, are those who enjoy reading, read faster, read more, understand most of what they read, and therefore enjoy reading. This is the corresponding 'virtuous circle'.

One can try to break out of the vicious circle and into the virtuous circle at any of the labelled points: one can try to read more, to read faster, to understand more (eg by enlarging one's vocabulary), or to enjoy reading more by finding more interesting or entertaining texts. Perhaps the best thing is to try to do all of

these things together. If there is one point in the circle which offers the most hope of making the breakthrough, it is the finding of more entertaining and motivating texts, but this is the one least amenable to investigation by controlled experiment.

Reading speed

Speed, by contrast, is highly amenable to experimental work. At the crudest level this can be carried out by assigning texts of known length, measuring the time taken, and administering comprehension questions to verify that some or all of the content has been understood. The times and scores can be plotted on graph paper so that progress during a course of training can be measured. This is the basic technique of many different reading speed courses, most notably the Reading Laboratories published since the 1950s by the Science Research Association of Chicago, where the individual texts are printed on separate colour-coded cards.

The problem with the Reading Laboratories, as with much published reading comprehension material, is that what is presented for reading is material that would normally be studied rather than read, ie short texts of factual content, monitored through comprehension questions demanding recall of detail. There is no doubt, however, that reading schemes like those of the SRA have been successful. They may have achieved their success by tapping in to the hunger that exists among learners for feedback. The elaborate record-books, with blank columns waiting to be filled with timings and scores, have had the effect of making children read more, even though what they were being asked to read was not intrinsically motivating.

Eye movements

Another approach to studying reading speed is to investigate the physical processes. This has been done using elaborate equipment, either helmets wired with electronic devices or by having readers wear special contact lenses which reflect light on to a sensitive surface, allowing eye movement to be registered with precision. Text can also be masked, so that the width of the perceptual window can be estimated. The research has been of great value in establishing some non-obvious and non-trivial facts about the physiological bases of reading.

The eye does not move steadily through a piece of text during the reading process; if it did, you could not read anything, any more than you can see things close beside the track when you are travelling on an express train. Instead, the eye moves in a series of hops, stopping to take in a chunk of text and then moving on to the next chunk. These hops are known as 'saccades' and the pauses between them as 'fixations'. What research has shown conclusively is that the rate of fixation varies much less than would be expected and that it is barely trainable. Adults and children, good and bad readers alike, make between three and four fixations per second.

Things are very different when we measure not fixations per unit of time but fixations per quantity of text. In a study by Taylor, Fracksenpohl and Pettee (1960), young readers were found to make 183 fixations per 100 words of text, while college-level readers made 75. Extrapolating these figures one arrives at reading speeds of 100 wpm for the youngest group (making 3 fixations per second) and a little over 300 wpm for the college level group (making 4 fixations per second). The figure of 300 wpm is a reasonable estimate of the average or 'cruising' speed of an adult reader, though with a great deal of variability due to the reader's purpose in reading and to the nature of the text itself. Various authorities have suggested that 180 wpm may be a threshold between immature and mature reading, that a speed below this is too slow for efficient comprehension or for the enjoyment of text. (See, for example, Bright and McGregor, 1970:96.) One question which is being investigated by the present authors is whether a cruising speed below 180 wpm is itself less flexible than higher speeds.

The perceptual window

The good reader who makes fewer fixations for a given amount of text appears to be using a wider perceptual window, ie to be recognising more characters per fixation, but the experimental evidence does not suggest that this will explain

differences in reading speed other than the gross differences between the youngest and the most mature readers. A spacing of 7 characters (about 1.3 words), which is roughly the separation between the fixations of college-level readers, corresponds to 2 degrees of visual angle at a normal reading distance, which in turn corresponds to the area of the retina called the fovea, ie that part where the cells are densest. Character recognition falls off sharply as the characters are seen by an area of the retina away from the fovea, and the eye cannot be trained to take in larger chunks of text at the purely physical level.

Guidance systems

There are two competing theories to explain how the brain chooses where to make fixations. Hochberg (1970) suggests that the visual system uses white space to guide the eye towards early word positions and across very short words which can be expected to have very little information content. This assumes that some information from the right periphery is being intelligently processed during each fixation in order to prepare for the next saccade. The competing view is that of Kolers (1966) who suggests that individual saccades are random, making no use of local information, though rhythmic and evenly spaced across larger segments. Spacing will be affected by the difficulty of the material but will settle into rhythmic regularity after a time. What experimental evidence there is slightly favours Hochberg's position, since it seems that the beginnings of long words get more than their statistically fair share of fixations. It may be that this is something which good readers do more effectively than weak readers. Certainly reading trainers often secure initial improvement by making readers run their fingers along the line of text. What this suggests is that one fault of the weak reader is making saccades which are so random that they miss the current line altogether.

Implications

If the main difference between the fast reader and the slow reader is neither the rate of fixation nor the breadth of the perceptual window, what can it be? The evidence suggests that both fast and slow readers see the same quantity of text per unit of time. Obviously fast readers are sampling the text at wider intervals, and are supplying the meaning of the unseen text between fixations by making predictions based on their familiarity with the language. Training the slow reader, in this case, must be a matter of providing enough familiarity with the language for guesses to be well grounded and of encouraging the reader to be adventurous enough to make some guesses.

Regarding the second of these, Bright and McGregor say "The English teacher can do half of this job in a few minutes by telling the pupils who are reading at about 100 words a minute 'You will now read better if you read faster.'" (1970:99) Pace Bright and McGregor, the problem is not so much one of telling as of persuading, since the assertion seems very counter-intuitive. Bad readers always believe sub-consciously that the way to understand more of a text is to concentrate on more of its component words. One way of persuading them otherwise is to give them larger chunks to look at for short times, and then demonstrating that something can still be understood even though every word has not been taken in. This is what the latest version of HOPALONG tries to do by providing two speed ranges, one using text broken up into short phrases, averaging 2.5 words, and one with longer phrases averaging 4.5 words. The main evidence of a gain in speed should come not from the changes in immediate speed as measured by the computer's clock, but rather from the pupil's own judgement that they can cope with the larger chunks as confidently as they can with the smaller.

Regressions

Most of the hops the eye makes are forwards, but virtually all readers occasionally make backwards hops, to check on something which has not been understood. These backwards hops are known as 'regressions'. Good readers make fewer regressions than bad readers, but the difference is not large enough to account for the overall difference in performance. Reading trainers used to make use of devices, known as tachistoscopes, which ran a fixed size window through a

text at a calculated speed. Clearly these devices make most regression impossible. They have fallen into disfavour as we have come to recognise that regression is part of the strategy of the skilled reader as well as of the weak reader. Useful information about the reader and the text can be obtained by seeing just when regressions occur and trying to find reasons for them within the text, but the cost of the equipment and the organisational problems of making it available rule it out except in experimental environments.

Computer equivalents of the tachistoscope exist. For example, there is an American training program called SPEED READ which flashes single phrases on screen at a fixed central position, or which flickers through a text displaying one phrase at a time. What is disturbing about such programs is that they are offered as training aids, and appear to be working on the false premise that reading speed can be increased by accelerating the number of fixations. The authors assume that their pupils will use the computer screen as the principal medium for reading during the training.

HOPALONG

HOPALONG is not intended as a training device, but simply as a research and diagnostic aid. This is a point that is repeatedly emphasised in the program documentation, and one of the authors' main worries is that the program will be misused by those who have ignored this warning. The aim of the program is to drive readers to paper for the bulk of their practice, and certainly not to make them depend on screen presentations.

HOPALONG derives from a prototype created on a BBC B microcomputer in 1987. Version 1 for MSDOS was completed in March 1988, and an improved Version 1.2 in September 88. Version 1.3, incorporating choice of phrase length and freedom to choose screen colours, has now been completed, but the bulk of the trials reported below were carried out using Version 1.2 running on a PC with a colour-graphics adaptor and colour monitor.

The program displays a text one screen page at a time and runs a highlight through it. Like the tachistoscope, it calculates a speed mechanically, but it differs significantly from the tachistoscope in that it does not mask out the page completely; instead it just prints the page in a light colour. It discourages regression but does not prevent it.

Colours and controls

In the original form of the program the text was printed in blue on black (quite legible on the screen of a BBC computer) while the highlighted phrase was in yellow. The current IBM version uses a white background in order to come closer to the normal experience of reading from paper, prints the text in green or cyan, and highlights the current phrase in black. All these colours can be changed at the time of loading, and the new colours saved as defaults.

The learner has control over the pace of the highlight. Touching the left arrow key slows it down while the right arrow key speeds it up, in increments of 10% in each case. Touching the down arrow halts the highlight completely; a second press of the down arrow re-starts it at the same speed. There are also keys to go back one page, go back to the beginning, or quit. The user is told to adjust the speed of the highlight until it feels comfortable, and is encouraged to use the controls, including the pause control, whenever they want to. We try to make it clear to the student that the program is not a race or a test of proficiency, simply a measure of comfort, and we also stress that what is being tested is not so much the student as the student's relationship to that particular piece of text.

The computer screen looks like the following illustration during the run of the program (in 40 column mode). At the start of the session the speed is set either to 120 wpm or to 180 wpm. The status line is constantly updated to show current speed and location within the full text, and it supplies a reminder of the most commonly used controls.

```
<Slower Page 1 120 wpm Faster>
          My chief duty
```

```
        as an office boy
      was to copy letters.
          Only very few
      of the communications
     were then typewritten.
           There were
        no female typists
    and only the most senior
          of the staff
   had male shorthand writers
        attached to them.
    All inter-office letters
      and ordinary letters
           to clients
          were written
         in copying ink
    by the clerks concerned,
      and all these letters
        had to be copied
           by damping
```

When the highlight reaches the foot of the page a new page scrolls in which includes the last phrase on the preceding page. A screen page contains about 220 words in 80-column mode or about 110 in 40-column mode, and the texts can be of any length from 2 to 25 pages. When the highlight reaches the end of the text, there is a short pause before the display switches to a graph showing the history of the reading. The graph can be copied to the printer, and the user also has the option of printing out the whole text on the printer. The printed text will include one of the following marks at each decision point:

[<]	slowed down
[>]	speeded up
[^]	went back one page
[^^]	went back to the beginning
[+]	came back to this point
[7]	pause of (eg) 7 seconds

This allows the learner to look at the text at leisure and try to see what, if anything, caused problems, or to take the text to a teacher, look at it together, and get advice.

The function of the highlight

The highlight itself is not intended to be matched to fixations or to guide them in any exact way. If it was, the length of phrase would need to be limited to 7 characters and the highlight would have to move between three and five times a second, causing an uncomfortable flicker. This effect does occur in Versions 1 and 1.2 of the program at the highest speeds available (650 wpm plus) and most users find it unpleasant or distressing. In version 1.3 the problem has been eliminated by linking the maximum speed to the average length of phrases, and offering two ranges, 50 to 300 wpm with short phrases and 100 to 500 wpm with long phrases. The highlight is gated so that it cannot move faster than two hops per second, and one would therefore expect the reader to make two fixations or more per exposure. The role of the highlight is very much that of a finger run through the text, guiding the eye towards the general zone where the next fixation can be made, and giving a rhythmic quality to the reading. We are interested in using the rhythm as evidence of speed and interruptions to the rhythm as evidence of problems in textual comprehension.

Reading from the computer screen

It is as well here to sound a warning about the physical characteristics of computer screens. The computer screen is not like a book page. The width is greater than the height; you often see light text on a dark background (though not in HOPALONG); the letters are usually much larger than book letters; the screen is

displayed in a vertical plane; and it is fixed in position, which means that you have to focus by moving your body rather than by moving the page. All of this makes reading a computer screen much more like reading notices on a board than reading a book. For these reasons the actual results of using this program will not be precisely the same as tests using ordinary print on paper. This applies particularly to good readers, who will probably not achieve as high speeds with the HOPALONG program as they could with an ordinary test. The range of speed which the program allows is from about 40 wpm, slow enough for anybody, up to about 600 wpm, which will be too fast for most of the target users. Even a good reader who can cope with print at 800 wpm will probably find that 600 wpm is the highest speed at which they feel comfortable on screen.

Modes

The program displays the texts either in 40-column or in 80-column mode. In 40-column each short phrase appears in large characters on a separate line, centralised on the page. The highlight moves down the page one line at a time. This is not a natural way of reading, but it is much easier for people who cannot read fast or who are having reading difficulties. In 80-column mode the text appears in smaller letters laid out like a conventional book page. The highlight moves along the line and then down to the beginning of the next line, imitating the normal movement of the eyes when we read a book. A 40-column screen contains about 70 words, while an 80-column screen contains about 220 words. In neither case is there as much text on view as there is with a normal book, where an open spread shows between 600 and 1000 words at a time.

Text divisions

The program makes use of text which is stored as plain ASCII files. The texts are broken up into phrases or 'chunked', a process which has to be carried out manually. The authors are designing a text editor which will carry out mechanical chunking, which can then be manually edited with less effort. If a purely mechanical method of chunking is used, we are bound to abandon the principle of preserving, as far as possible, meaningful phrase boundaries, since there is no way one could build enough artificial intelligence into the chunking routines, though it might be possible to include a few rules such as "do not allow a or the as the last element of a chunk". The loss of intelligent chunking might turn out to be unacceptable if the improper divisions affect the mental chunking that readers obviously carry out in the comprehension process. Carver (1970) demonstrated that division of a text into meaningful phrases led to no improvement in reading over text presented in narrow newspaper columns, which suggests that visual chunking is unimportant. However highlighting and arbitrary chopping into columns may not be totally comparable as chunking methods. Only further experimental work will tell.

Comprehension questions

The program is supplied with a facility to present multiple choice questions from a linked text file. Simply knowing that there are questions to be answered may affect the way that students read, but the questions themselves are subject to all the ordinary drawbacks of comprehension questions and are an extremely fuzzy measure of comprehension. We included the questions as much to meet teachers' and students' expectations as to gain information, and we play down the value of the questions both in the face-to-face instructions we give and in the program documentation.

At the beginning of the program the student sees an introductory page which says what the topic of the text is and presents two questions to think about during the reading, one factually related to the content, such as

> Did Mr Lee have a happy marriage?

and the other a personal reaction question such as

> Do you think office boys nowadays find their work
> easier or more interesting than the writer's?

These questions are never answered by the program itself, of course, but are intended to provide material for follow-up discussion in a feedback session.

Participants

The subjects in our trials have all been overseas students at Bristol University. These included some attending special English courses, with entry levels from Band 3.5 to Band 5.5 on the ELTS test, and some who were admitted to graduate study with entry levels of 6.0 to 7.5. A number of native-speaker colleagues have also tried out the program and reported their reactions.

Although HOPALONG is in no sense designed as a training device, it is nevertheless the case that you have to get to know it in order to use it effectively. Otherwise unfamiliarity with the controls is likely to get in the way of the reading. Our most interesting results have come from a group of in-house students who were able to use the program on four or more occasions, while results from external students who came in for one session only were clearly unreliable. One student, for instance, at his first attempt accelerated to a high speed, made few adjustments after that, and then scored zero on the comprehension. No doubt he would have trained himself given the opportunity.

Results

Results of these trials appeared in an M Ed dissertation presented in 1989. As one would expect, a gross correlation emerged between the learner's standard of English, the difficulty of the text, and the chosen speeds, but the curves do not match smoothly. The discriminatory power and usefulness of the graphs diminishes when text turns out to be far too easy or far too difficult. To explain the irregularities one would have to run separate trials with doctored text to test out each hypothesis.

It was also clear that reaction to a reading difficulty, whether in vocabulary or syntax, was not immediate enough to make identification of the problem easy when studying the graph or the text printout. It seemed rather that the decision point showed where an accumulation of previous difficulties had created a crisis of understanding. This might, of course, change when users become more familiar with the controls, especially the pause control which was very little used by the overseas students.

Readers varied widely in their strategies, showing various degrees of adventurousness. Three students read with no change of speed, while one made 133 adjustments within a single short journalistic text, though the same reader made only 22 adjustments on the presumably more familiar science text. One observed tendency was for more adjustments to be made by the slow readers than the fast ones. More changes were made near the beginning of text, suggesting that readers did eventually settle to a cruising speed.

What you see

The work we have done does no more than support what is already well known, namely that the good reader, in particular the fast reader who achieves speeds well over 300 wpm, is actually seeing less of the text that the bad reader, but is able to supply the missing information from expectations about grammar and collocation, ie from general familiarity with the language. The consequence seems to be that one cannot become a good reader without having the competence to make predictions. That in turn means one should be doing one's reading with material within which one can make predictions, ie material which has been simplified appropriately or which deals with subject matter which is conceptually familiar.

Neither HOPALONG nor any other dedicated reading improvement scheme is likely to make much difference to predictive skills. All that it can do is provide an additional means for learners (and to some extent their teachers) to find out more about themselves and about the text they are trying to read. It invites learners to think about speed and their feelings about a piece of text. That may be all that is necessary for them to take charge of their own learning and work towards an improvement.

References:

Bright, John and McGregor, Gordon (1970). Teaching English as a Second Language. London, Longman.

Carver, R.P. (1970). "Effects of a 'chunked' typograph on reading with

comprehension." Journal of Applied Psychology.
Hochberg, J. (1970). "Components of literacy: speculations and exploratory research". In Levin, Harry and Williams, Joanna (eds), Basic Studies on Reading. New York, Basic Books.
Kolers, D.A. (1966). "Reading and talking bilingually". American Journal of Psychology, 99, 357 - 376.
Nuttall, Christine (1982). Teaching reading skills in a foreign language. London, Heinemann Educational Books.
Science Research Associates (1959 onwards). SRA Reading Laboratories. Chicago, Science Research Associates.
Taylor, Fracksenpohl and Pettee (1960). Grade level norms for the components of the fundamental reading skill. Huntington, N.Y. Educational Developmental Laboratories Inc. Bulletin 3/1960.