

Extracting Text from PostScript

craig g. nevill-manning, todd reed and ian h. witten
Computer Science Department, University of Waikato, Hamilton, New Zealand
(email: {cgn,treed,ihw}@cs.waikato.ac.nz)

SUMMARY

We show how to extract plain text from PostScript files. A textual scan is inadequate because PostScript interpreters can generate characters on the page that do not appear in the source file. Furthermore, word and line breaks are implicit in the graphical rendition, and must be inferred from the positioning of word fragments. We present a robust technique for extracting text and recognizing words and paragraphs. The method uses a standard PostScript interpreter but redefines several PostScript operators, and simple heuristics are employed to locate word and line breaks. The scheme has been used to create a full-text index, and plain-text versions, of 40,000 technical reports (34 Gbytes of PostScript). Other text-extraction systems are reviewed: none offer the same combination of robustness and simplicity. © 1998 John Wiley & Sons, Ltd.

key words: PostScript interpreter; text extraction; digital libraries; full-text indexing; page layout

INTRODUCTION

It is often useful to be able to extract plain text from PostScript files. For example, the New Zealand Digital Library* contains a collection of computer science research reports that is full-text indexed; to create this index it is necessary to extract all the text from the document files that constitute the originals of the reports. Similar requirements for recovering plain text from the page description arise in many applications.

Converting PostScript to plain text is like an optical character recognition process for which the input is not 'optical' but rather an electronic description of a page. That makes the problem sound trivial, for of course the characters need not be 'recognized' because the document file represents them symbolically, as ASCII codes, rather than pictorially as ink on paper. However, it is complicated by the fact that PostScript is more than a page description language; it is a programming language. Identifying the textual content involves more than conversion from one format to another – it involves predicting the output of a program.

Moreover, although sometimes it is just the stream of words that needs to be extracted (for example, for indexing purposes), other applications will need an approximation to the formatting information so that the text can be reproduced in different ways. A common requirement is to produce a (necessarily) approximate

* <http://www.nzdl.org>

representation of a PostScript document in HTML, the hypertext mark-up language used in the World-Wide Web, and for this, structural features such as paragraph boundaries and font characteristics must be identified.

We describe a simple and robust scheme for extracting text from PostScript. Essentially, it involves modifying the behavior of a PostScript document so that, when executed, it produces ASCII text rather than pixels on a page. We also describe heuristics that improve the detection of word and line divisions by employing global information. Because of differences in character spacing, text with hard-coded line breaks that correspond to the layout of lines on the page exhibits grotesquely ragged right margins when displayed on a screen with a different font. Consequently, we have developed simple heuristics to recognize paragraph boundaries, and to solve other practical problems with the output of the extraction process. However, we make no real attempt to reconstitute a good approximation to the original page layout and formatting; that lies beyond the scope of this article.

These techniques have been tested and proven to be robust on a collection of over 32,000 documents from a wide variety of sources (300 FTP sites) that constitute the computer science research collection in the above-mentioned digital library. Other public-domain PostScript extraction programs proved to be either very unreliable, or far too computation-intensive for our purposes. It is worth noting that our scheme is robust enough to cope with Adobe Portable Document Format (PDF) files* without any modification at all.

Extracting text from PostScript is challenging both from a fundamental computational perspective and from a pragmatic programming stance; the next section explains why. We then describe how a small PostScript program can perform the basic task elegantly and reliably. However, it makes simplistic assumptions about line and character spacing, so we present improvements to the basic algorithm which identify text boundaries more robustly and also solve some other practical problems that arise. The final section compares this technique to others and situates it in the spectrum of what is available in the public domain.

THE PROBLEM WITH POSTSCRIPT...

The programming language PostScript is designed to convey a description of a desired page to a printer.¹ Extracting text from PostScript files poses problems at two quite different levels. The first is at an abstract level. Because PostScript is a programming language, its output cannot, in principle, be determined by a textual scan. The second problem is a more pragmatic one. Although almost all actual PostScript documents have a fairly straightforward computational structure that makes the output text immediately available, it is fragmented and inextricably mixed with other text that does not form part of the desired output.

A problem in principle

In principle, the process of extracting text from a PostScript document is not one of format conversion, but rather a matter of predicting the output of a computer

* See www.adobe.com for details.

program. Consider the PostScript sample and output in [Figure 1](#). This admittedly rather contrived example is a program that displays the sixth number of the Fibonacci sequence 1, 1, 2, 3, 5, 8, namely '8'. However, the character '8' appears nowhere in the input document: it is calculated recursively by the *fib* function. Any syntactic approach to conversion will be unable to accommodate such input and, to the extent that it occurs in practice, will be unreliable.

The only way to know the output of a program is to execute it. Thus, any robust scheme for extracting text from PostScript must execute the document using an appropriate interpreter. Of course, the execution of the document must be altered to allow the text that is produced to be captured instead of rendered by placing pixels on the page. This paper describes just such a scheme.

A problem in practice

Most PostScript files do not contain code as subtle as the recursive function of [Figure 1](#), and this may be regarded as a pathological case. [Figure 2](#) shows excerpts from a more representative example of a document file, along with the text extracted from it. Characters to be placed on the page invariably appear as parenthesized strings in the source.

An obvious approach to text extraction, used in some conversion programs, is to extract and concatenate all such strings. There are two problems with this simple strategy. First, font names, file names, and any other internal string information is represented in the same way in the file; examples can be seen in the first few lines of [Figure 2](#). Second, the division of text into words is not immediately apparent: spaces are usually implicit in the positioning of characters on the page. Text is often written out in fragments, so that each parenthetical string may only represent part of a word. Deciding which fragments should be concatenated into words is a difficult

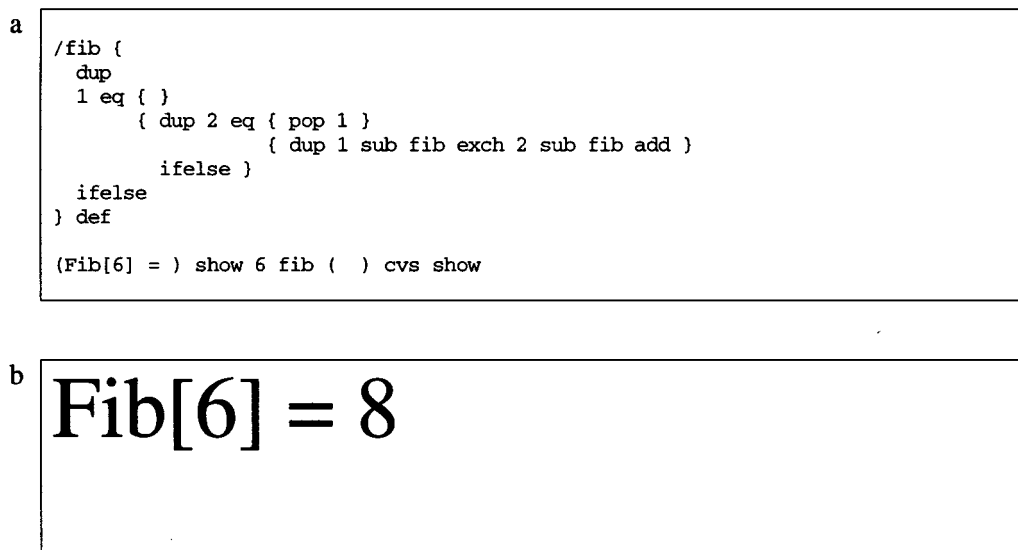


Figure 1. A PostScript program that prints a character that does not occur in the text of the program. (a) Program, (b) rendered output

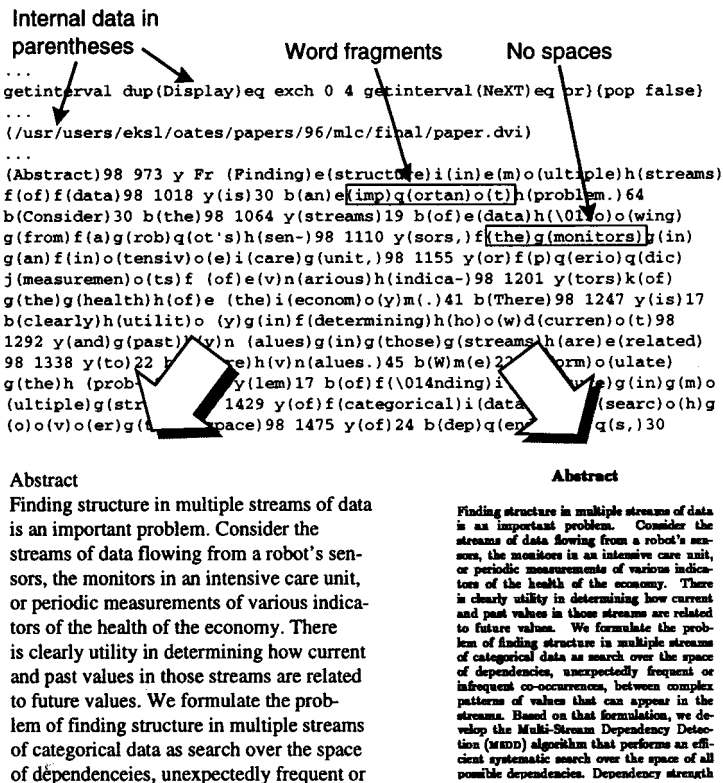


Figure 2. A PostScript document and the text extracted from it

problem. Although simple heuristics might be devised to cover most common cases, this is unlikely to lead to a robust solution that can deal satisfactorily with the wide variety of files found in practice.

Another drawback to the general approach of analyzing parenthesized strings in PostScript files is that it requires modification to handle PDF files, which are compressed, and contain no readable text.

A SIMPLE TEXT EXTRACTOR

Our approach is to modify a document by prepending a small prefix which, when executed by a standard PostScript interpreter, has the effect of producing ASCII text in a file rather than pixels on a page. To execute the code, the modified document can be processed either by a previewer such as *ghostscript*, or by sending it to a printer. Details of the procedure are given in the Appendix.

The trick is to redefine the PostScript *show* operator, which is responsible for placing text on the page. Regardless of how a program is constructed, any text that appears on a page passes through this operator (or one of the five variants discussed below). We redefine *show* to write its argument, a text string, to a file instead of rendering it on the screen. When the document is executed a text file is produced instead of the usual set of physical pages.

We explain the operation by incrementally improving a simple program. Prepending the line `/show { print } def`, shown in Figure 3(a), to the document of Figure 2 redefines the `show` operator to print to standard output. The result is shown at the right of Figure 3(a). One problem has been solved: separating what text is destined for a page from the remainder of the parenthesized text in the input file.

The question of identifying whole words from fragments must still be addressed. The text in Figure 3(a) contains no spaces, and needs to be segmented into words. Printing a space between each fragment results in the text in Figure 3(b). Spaces do appear between each word, but they also appear within words, such as *multiple* and *important*.

To detect where spaces should be inserted, it is necessary to consider how

a	<pre>/show { print } def</pre>	<p>Finding structure in multiple streams of data is an important problem. Consider the streams of data showing from a robot's sensors, the monitors in an intensive care unit, or periodic measurements of various indicators of the health of the economy. There is clearly utility in determining how current and past values in those streams are related to future values</p>
b	<pre>/show { print () print } def</pre>	<p>Finding structure in multiple streams of data is an important problem. Consider the streams of data showing from a robot's sensors, the monitors in an intensive care unit, or periodic measurements of various indicators of the health of the economy. There is clearly utility in determining how current and past values in those streams are related to future values</p>
c	<pre>/X 0 def /show { currentpoint pop X sub 5 gt { () print } if dup print systemdict /show get exec currentpoint pop /X exch def } def</pre>	<p>Finding structure in multiple streams of data is an important problem. Consider the streams of data showing from a robot's sensors, the monitors in an intensive care unit, or periodic measurements of various indicators of the health of the economy. There is clearly utility in determining how current and past values in those streams are related to future values.</p>
d	<pre>/X 0 def /protoshow { currentpoint pop X sub 5 gt { () print } if dup print systemdict exch get exec currentpoint pop /X exch def } def /show {/show protoshow} def /kshow {/kshow protoshow} def /widthshow {/widthshow protoshow} def /ashow {/ashow protoshow} def /awidthshow {/awidthshow protoshow} def</pre>	

Figure 3. Simple and complex ways of producing text from PostScript. (a) Printing all fragments rendered by `show`; (b) putting spaces between fragments; (c) putting spaces between fragments more than five points apart; (d) extension to other PostScript `show` operators

fragments are placed on the page. Figure 4 plots the distribution of distance between fragments, measured in printers' points,* in one technical report. The peak around zero corresponds to the fragments being part of the same word, while the broader hump around 16 corresponds to the inter-word spacing. The latter is spread out by the variation that occurs in word spacing. This is particularly striking when the margins are 'justified' by adding extra space to pad out short lines. It seems reasonable from the graph to pick a threshold of about five points, below which fragments will be concatenated and above which spaces will be inserted.

The prologue in Figure 3(c) implements this modification. The variable *X* records the horizontal co-ordinate of the right-hand side of the previous fragment. The new *show* procedure obtains the current *x* coordinate using the *currentpoint* operator (pop discards the *y* coordinate), and subtracts the previous coordinate held in *X*. If the difference exceeds a preset threshold, in this case five points, a space is printed. Then the fragment itself is output.

To record the new *x* coordinate, the fragment must actually be rendered. The line `systemdict /show get exec` retrieves the original definition of *show* from the system dictionary and executes it with the original string as argument. This renders the text and updates the current point, which is recorded in *X* on the next line. The execution of the original *show* operator is one of the innovations of this technique over previous ones: it provides a simple and foolproof way of updating coordinates. This new procedure produces the text at the right of Figure 3(c), where all words are segmented correctly. Line breaks can be detected by analyzing vertical coordinates in the same way and comparing the difference with a fixed threshold. There is one remaining problem in the text of Figure 3(c) – a *fl* ligature is represented by the symbol – and we return to this in the next section.

The four variants of the PostScript *show* command, namely *ashow*, *widthshow*, *awidthshow* and *kshow*, are treated similarly. A procedure is defined to do the work. It is called with two arguments, the text string and the name of the appropriate *show* variant. Just before it returns, the code for the appropriate command is located in the system dictionary and executed. Figure 3(d) shows the complete prologue.




Figure 4. Distribution of inter-fragment distances

* There are 72 points to the inch.

Notwithstanding the use of fixed thresholds for word and line breaks, this scheme is surprisingly effective on the 40,000 technical reports in the technical report repository. In the next section, we describe several enhancements that further improve performance.

AN IMPROVED TEXT EXTRACTOR

Although the solution developed in the previous section is elegant and robust, it has several minor problems that can be solved by introducing some heuristics. These problems include incorrect word-boundary identification when using very large or very small fonts, the need to distinguish paragraph and line breaks, dealing with non-ASCII characters, dehyphenation, and page reversal. We address each in turn. They are implemented in the Python language³ as a post-processor to the PostScript scheme just described. Because it is necessary to have access to detailed information on word and line spacing, a simple modification is made to the PostScript prologue in [Figure 3\(d\)](#) to provide this information.

Spacing in large and small fonts

The use of a fixed threshold to distinguish inter-word gaps from inter-fragment ones fails when the text is being printed in a very large or very small font. When using large fonts, as in titles, inter-fragment gaps are mistakenly identified as inter-word gaps, and words are broken up into individual letters. When using small fonts, as in footnotes, inter-word gaps are mistaken for inter-fragment gaps, and words are run together.

To solve this problem, we express the word break threshold as a fraction of the average character width. This width is calculated for the fragments on each side of the break by dividing the rendered width of the fragment by the number of characters in it. Examination of histograms of breaks as a fraction of character width leads to a threshold of 30 per cent, which is effective for a wide range of type sizes. This modification also eliminates all dependence on the units in which the co-ordinate system is expressed.

Paragraph vs. line breaks

The line breaks in a PostScript document are designed for typeset text. Plain text does not always share the same line width, and it is often desirable to wrap lines differently depending on the situation—for example, the width of the window containing the text. Paragraph breaks, on the other hand, have significance for the document's content and should be preserved.

Paragraph breaks can be distinguished in two ways. The first presupposes that more space is left between paragraphs than between lines, as in [Figure 5\(a\)](#). In this

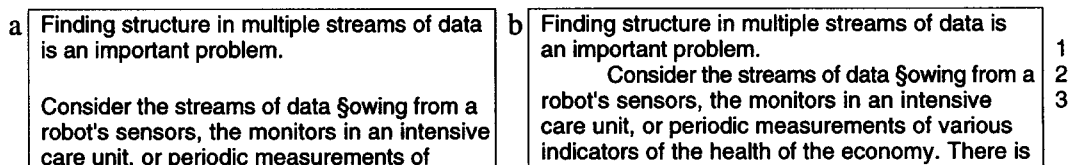


Figure 5. Two styles for paragraph breaks. (a) Additional space; (b) indentation

case, any break exceeding the average line space can be treated as a paragraph break. For ‘average line break’ we use the most common non-trivial change in y coordinate throughout the document.

The second way separates paragraphs marked by indentation rather than line spacing. In [Figure 5\(b\)](#), the first line of the second paragraph is indented from the left margin. This is often sufficient for identifying new paragraphs, but additional heuristics can reduce the number of mis-identifications. Consider the three lines numbered to the right of [Figure 5\(b\)](#): one at the end of a paragraph and the next two beginning the subsequent paragraph. For these to signal a paragraph break,

- (a) the second line must be indented relative to the other two;
- (b) the second line must start with a capital letter;
- (c) the third line must be longer than the other two;
- (d) the second line must be almost as long as the third.

In the last two constraints it is the text length that is measured, not the distance between margins. The first line terminates a paragraph, and will not usually extend to the right margin, while the second one is indented. While certainly not infallible, these rules work well in practice.

Non-ASCII characters

When producing ASCII text from PostScript some character translation is inevitable because the character sets used in documents are much richer than the ASCII set. In technical documents, the most common use of unusual glyphs is in mathematical formulas. These include some ASCII characters – for example, variable names and digits – and some other symbols – Greek characters, integral signs and the like. The latter must either be approximated in ASCII, or deleted. We decided to flag unknown characters with a question mark because there can be no truly satisfactory character representation of mathematical formulae.

Other commonly-used glyphs include ligatures, bullets, and printer’s quotes. Ligatures are specially joined and kerned combinations of characters such as *fi* and *fl* that are rendered as a single glyph. This occurs very frequently in \TeX output – indeed, an example appears in [Figure 3](#) – and since this is common in our environment, we recognize such symbols and map them to their two-character equivalent. Bullets and printers’ quotes (“” rather than ’’) are recognized in the same manner. A general solution to this problem requires recognizing of particular font encodings and mapping glyphs to meaningful ASCII representations. We investigated this problem, but were unable to devise a universally effective mechanism for detecting encoding changes.

Dehyphenation

When documents are justified to a fixed right margin, words are often hyphenated. We reverse this process by detecting hyphens at the end of a line and concatenating the fragments on either side of the break. Admittedly this can remove the hyphen from a compound word that straddles a line break, but such situations are rare. Dehyphenation does not work across page breaks because headers and footers separate the last line of one page from the first line of the next.

Page reversal

The pages of PostScript documents often appear in reverse order. This is for mechanical convenience: when printers put pages face up on the output tray, the first page produced is the last page in the document. The Adobe document structure² defines a convention for specifying page ordering, but it is rarely followed in actual document files.

To detect page order, several heuristics were considered but rejected on grounds of unreliability. Our final solution is to extract numbers from the text adjacent to page breaks. These are usually page numbers, and one can tell if a document is reversed by checking whether the sequence is increasing or decreasing. This is fairly reliable because even though some numbers in the text are erroneously identified as page numbers, the decision is made on a global, majority, basis.

OTHER SCHEMES

This system was developed as a solution to a practical problem that other systems failed to solve. Here we briefly describe similar schemes. The first two attempt extraction by a textual scan, while the remainder make use of a PostScript interpreter.

<code>ps2ascii.pl</code>	This is a widely-used Perl script ⁴ that extracts parenthesized text from a PostScript file. <code>ftp://ftp.cs.brown.edu/pub/postscript/ps2ascii.pl</code>
<code>ps2txt</code>	This is a stand-alone C program that extracts parenthesized text from a PostScript file. It also includes some special code to deal with files generated by the <i>dvips</i> program. <code>ftp://ftp.funet.fi/pub/archive/alt.sources/volume92/Feb/920223.01.gz</code>
<code>ps2a.sh</code>	Written entirely in PostScript, this system requires <i>psh</i> , a component of the NeWS windowing system – although it could run using <i>ghostscript</i> or a printer. It is a rather complex program, and is optimized to perform with output from <i>t_EX</i> . Like the system described in the present paper, it does execute the original <i>show</i> operators in the PostScript input. <code>ftp://macadam.mpce.mq.edu.au/pub/comp/src/ps2a.sh</code>
<code>pstotext</code>	This system, from Digital Equipment Corporation, includes components in both PostScript and C. Although very comprehensive, is it correspondingly complex, and extremely slow to execute. It succeeds in tracking font changes and does an excellent job of conversion. <code>http://www.research.digital.com/SRC/virtualpaper/pstotext.html</code>
<code>ps2ascii</code>	This program, which is included with the standard <i>ghostscript</i> distribution, is complex but, in our experi-

ence, not at all robust.
<http://www.cs.wisc.edu/~ghost>

ps2ascii [ucf] This variant of the *ghostscript* **ps2ascii** was developed at the University of Central Florida. It includes components in both PostScript and Perl. (no longer available)

ps2html This variant of **ps2ascii** was developed at Johns Hopkins University to convert journal articles to HTML. It strives to preserve the formatting of the original PostScript document. However, it is designed exclusively for documents with a format specific to certain journals of the Johns Hopkins University Press, and recognizes – and gives special attention to – PostScript files that were generated with a specific software package (QuarkXPress). A table containing numerous parameters is used to aid conversion, and can, in principle, be modified for new formats.
<ftp://bradley.bradley.edu/pub/guru/ps2html/home.html>

Much research has been conducted on the general problem of analyzing document images. Although most of this is concerned with the higher-level logical structure of documents, some schemes look at the problem of distinguishing line from paragraph breaks.^{5,6} These schemes are considerably more complex than the simple heuristics described here, the main thrust of our work being on the problem of extracting plain text files from PostScript. It would be interesting to see if they gave significant improvement on a large collection of documents.

CONCLUSIONS

We have presented a simple, elegant and robust scheme for extracting text from PostScript. The heuristics it uses to reconstitute word spacing and paragraphing, and to restore other small aspects of the text, work well although they are obviously not infallible. We have not yet performed a quantitative evaluation of these heuristics, partly because the notion of a random sample of documents is not well-defined – everything depends on how the PostScript is generated in the first place, and there is no standard corpus of example files that are used to evaluate text conversion schemes. However, a large-scale sample from a wide variety of sources can be viewed in the New Zealand Digital Library collection of computer science research reports.*

In future, we will address the extraction of further structural information from layout, ranging from section headings to bibliography entries. We anticipate that deficiencies will inevitably continue to emerge as new documents, with new layout styles, are encountered, and we will have to address them by evolving the system constantly. To prevent it developing into a large, unstructured tangle of heuristic hacks, which will rapidly become unmaintainable, we propose to investigate the application of machine learning techniques to infer a rule set from example reports.

* <http://www.nzdl.org>

Whenever a new format is encountered, we would augment the set of training examples to include it and relearn the set of rules from scratch. Meanwhile, our heuristic system already does a good job of extracting text and rudimentary structure from PostScript documents.

APPENDIX: OBTAINING THE EXTRACTION SYSTEM

The simplest way to employ this scheme is to copy the prologue from [Figure 3\(d\)](#), prepend it to a PostScript file by hand, and run it through a PostScript interpreter. *Ghostscript* is an excellent public-domain interpreter, and is available from www.cs.wisc.edu/~ghost. Once it is installed, save the prologue above in a file called *prescript.ps*, and type

```
% gs -q -dNODISPLAY -soutfile=outfile prescript.ps infile.ps -q
```

where *infile.ps* is the PostScript file, and the text is to go into *outfile*.

An intriguing option exists on a Macintosh, or any other system that provides a mechanism for receiving data from a direct connection to a PostScript printer: the interpreter in the printer itself can be used. On the Macintosh, use BBEdit (www.barebones.com) to open the PostScript file, paste in the prologue, send the file to the printer, and capture the returned text. Because the *showpage* operator is redefined, the printer will not output any pages.

To take advantage of the improvements to the basic algorithm, download the *prescript* system from the 'technology' link on the New Zealand Digital Library page (www.nzdl.org). This consists of an augmented prologue, *prescript.ps*, and a Python script, *prescript.py*, that implements the heuristics described above.

REFERENCES

1. Adobe Systems Incorporated, *PostScript Language Reference Manual*, Addison-Wesley, Reading, MA, 1985.
2. Adobe Systems Incorporated, *Document Structuring Conventions*, Addison-Wesley, Reading, MA, 1989.
3. M. Lutz, *Programming Python*, O'Reilly & Associates, Sebastopol, CA, 1996.
4. L. Wall, T. Christiansen and R. L. Schwartz, *Programming Perl*, O'Reilly & Associates, Sebastopol, CA, 1996.
5. S. Tsujimoto and H. Asada, 'Major components of a complete text reading system', *Proceedings of the IEEE*, July, 1133–1149 (1992).
6. D. Rus and K. Summers, 'Using non-textual cues for electronic document browsing', in N. Adam, B. Bhargava and Y. Yesha (eds.), *Digital Libraries, Current Issues: Lecture Notes in Computer Science 916*, 129–163, Springer-Verlag, 1995.