Springer Link

Search

Conference of the Canadian Society for Computational Studies of Intelligence
AI 2006: Advances in Artificial Intelligence pp 300-311 | Cite as

Using Natural Language Processing to Assist the Visually Handicapped in Writing Compositions

Authors | Authors and affiliations

Jacques Chelin, Leila Kosseim, T. Radhakrishnan

Part of the Lecture Notes in Computer Science book series (LNCS, volume 4013)

Abstract

Using Natural Language Processing to Assist the Visually Handicapped in Writing Compositions

Jacques Chelin, Leila Kosseim, and T. Radhakrishnan

Department of Computer Science and Software Engineering,
Concordia University, Montreal, Canada
{jj_chel | kosseim | krishnan}@cse.concordia.ca

Abstract. Over the last decades, more and more visually handicapped students have attempted post-secondary studies. This situation has created many new challenges. One of them is the need to study text and electronic documents in depth and in a reasonable time. Blind students cannot flip through the pages of a book, skim through the text or use a highlighter. In this paper, we propose a solution in the form of an experimental prototype and show how natural language processing techniques can profitably assist blind students in meeting their academic objectives. The techniques used include the automatic creation of indices, passage retrieval and the use of WordNet for query rewriting. The paper presents a technology application of a practically usable software.
The system was evaluated quantitatively and qualitatively. The evaluation is very encouraging and supports further investigation.

1 Introduction

The visually handicapped have consistently progressed over the last decades in their efforts towards inclusion in the mainstream[7]. Integration in education and professional life in particular was possible due to the deployment of computers in every day life, without which it would not have been possible, or at least not to the same degree. In the wake of this integration, more and more blind students are attempting post-secondary studies.

This situation has created many new challenges and new needs specifically related to the in depth study of documents in a reasonable time so as to produce assignment submissions and research papers. In this paper, we discuss how Natural Language Processing (NLP) techniques can assist blind students in meeting their academic objectives. We present a prototype system built as an information probing and gathering environment. Its goal is to reduce the time it takes a student to do research on a specific topic and ultimately produce a paper in a time frame close to their sighted friends.

1.1 The Problem

On a regular basis, post-secondary students, especially those in Liberal Arts, have to do research on specific topics. Their task consists in consulting a wide variety of books, documents and Web sites and producing an essay, anything from a few pages to a full-fledged thesis. To be able to better understand the difficulties that visually handicapped students face when trying to write a paper, we need to give a quick description of the tools they use to read and write.

The visual handicap can be divided into two broad categories, the partially sighted and the totally blind. We only address the latter here. One way to make up for the absence of sight is the use of speech. The way computer technology is used in this case is to provide spoken output as screen readers that use an internal speech synthesizer. Words appearing on the screen are read aloud. The major problem with using speech as a medium is that it provides a very small working window due to the constraints of short term memory and linearity of speech. [6] provided evidence that people can remember about 7 chunks (in our case, terms) in short-term memory. [3] goes even further and suggested as little as 3 to 5 terms.

The other way to compensate for the loss of visual ability is through the sense of touch. For two centuries now, the blind have been able to read by moving their fingers across raised dots on thick paper. The appearance of computer based 'Braille displays' in which each dot can be raised or lowered through electromechanical devices has caused a huge leap in the accessibility to electronic documents. The major constraint here is that the user can only 'see' a 40-character window.

1.2 A Comparison

To write a paper, a student normally scans a substantial amount of documents quickly and easily using fast reading techniques; develops an outline while reading one or more of these documents more systematically, highlighting and taking notes; reviews and rearranges his notes and uses them to create a draft of his document; refines the document iteratively until it is complete; at all times, refers back to previous readings and versions.

At first sight, all these steps may seem easy, even mundane. However, for the visually handicapped, several problems exist. To name only a few:

- The amount of material to be read is, by itself, often a challenge – even for a sighted student.
- Reading difficulties go from very mild to very severe for handicapped students. This can be due to many factors such as a lack of reading and summarization skills, visual handicaps, dyslexia, ...
- Note taking is an art in itself. The classic index card method facilitates sorting of notes but is tedious and can hardly be used by students with visual handicaps.
- Contrary to sighted users who can see a full screen of information, the blind have no overview of documents. Everything is always seen through a 40 character window for Braille users, or a window of less than 8 spoken words for speech users.

What is more, the above problems are encountered while performing all major steps in writing compositions: Research, Analysis, Outlining and Composition.

2 Previous Work

As witnessed by the CSUN series of conferences (e.g. [1, 2]), the topic of hardware and software applications to help the visually disabled has received a lot of attention. However, the resulting software applications are either geared at reading or writing, but not both. In addition, most of the features provided are text annotation tools and very little NLP techniques are used. As many of these systems are commercially available, no paper describing their inner working seems to be available.

WYNN[1], Kurzweil 1000[2] and textHELP[3] are all tools to read and create documents mainly targeted for learning disabled students or users with learning difficulties like dyslexia, attention deficit disorder (ADD) and other literacy difficulties. These products typically provide text annotation tools such as bookmarking, note taking or outlining facilities. However, these facilities are often crude. For example, the user often cannot directly go to the position of a bookmark in the original document (through a hyperlink for example), or integrate bookmarks into existing outlines; thus limiting his access.

To our knowledge, NLP techniques used to improve the reading and writing tasks of the visually handicapped include only word prediction (in textHELP and WYNN) and homonym checking (in textHELP). Homonym support provides auditory and visual reinforcement of commonly confused like-sounding words. To avoid the confusion between homophones, the program color codes confusable words and lists possible alternatives with audible definitions and sample sentences. Word prediction allows the application to predict the most likely word to be typed given the previous context. The user types a letter and the program offers a list of the most likely words beginning with that letter. If the required word is on the list, it can be quickly selected. If the word is not on the list, typing the next letter will bring up a different choice and so on. Again, as these systems are commercial, it is not clear if a language model is used, or if a simple dictionary look-up is performed.

A related research project is that of [9], who developed an authoring environment to assist users in writing hypermedia documents. The system, based on a cognitive model of writing, offers features such as an outliner to help organise ideas and for creating and manipulating view areas. However, as this tool was not designed for the visually impaired, these features do not specifically address their needs. In addition, it provides no feature for skimming documents or other features to help find the content to be presented in the final document.

3 Proposed Solution

In order to assist the visually handicapped, a prototype system called ESCAS[4] was developed. The system can be seen as an information probing and gathering

[1] `http://www.freedomscientific.com/LSG/products/wynn.asp`
[2] `http://www.kurzweiledu.com`
[3] `http://www.texthelp.com`
[4]

environment or more practically as a text editor with features to assist in reading and writing documents. Its goal is to reduce the time it takes a student to do research on a specific topic and ultimately produce a paper in a time frame close to their sighted friends. To achieve this, the system assists the user in:

- Determining the relevance of documents without having to read them entirely.
- Accessing information related to a user specific theme or subject much faster than traditional methods.
- Freeing the user from tedious searches and reading chores for more productive and creative work.

The purpose here is to investigate where and how far NLP techniques can facilitate the access of students to information pertinent to their research. These techniques will convey additional information about the content of a document and faster access to relevant positions within the document.

Text annotation tools, such as the facility to create and annotate outlines, bookmarks, notes ... are the typical tools offered to help the visually handicapped (see section 2). These tools facilitate the access and the composition of documents, but cannot help in manipulating their *content*. We therefore looked at various NLP techniques to zoom in on material within an input document that is particularly relevant to the subject being developed by the user. In the case of sighted users, this can be achieved by skimming or fast reading through the material. Our intent here is not only to replace skimming with what NLP can offer but also go beyond skimming and implement such functionality that could be of help even to those who have no handicaps.

3.1 Indexing

When skimming literature, the user is looking for specific information quickly. To achieve this, a useful support is an index. In a standard index at the end of a paper book, important terms (words or phrases) are arranged alphabetically and are associated with the page numbers where they can be found. Blind students generally use scanned versions of paper documents. As the books were designed for paper medium, they do not contain hyperlinks to help users navigate quickly. The scanned version of an paper-book index is therefore of no help to blind users. They need direct access to the actual segment of text (be it a paragraph, a sentence ...).

Many approaches to index generation have been proposed to find terms that are representative of the document and to relate them to each other semantically (e.g. [4], [10]). In ESCAS, our goal is to generate many automatic pointers to text excerpts and to perform this task on the fly whenever a new document is loaded in the system. A syntactic analysis or chunking of the document is therefore not an option as it may not be fast enough and may not be able to parse all sentences. As we are more interested in recall than in precision, we opted for a non-linguistic approach.

In ESCAS, each index term is considered as a text node. Visually, under that node, instead of a page number, the user can see the actual text excerpt which contains the entry. By just pressing a key (*Enter* in our case) on an occurrence, the user is taken to its exact location in the text. For the user, the index is handy because, in addition to giving access to the original text segment, the index can be used to scan through a list of sentences rather than to have to keep switching to the text from a list of page numbers. The index includes both single-word terms and two-word terms. First, the text is tokenised, stopwords are removed and the remaining words are stemmed using the Porter stemmer [8]. For each stem, we then keep a list of pointers to the paragraphs in which they occur and to the position of the stem in the text. We also keep one of the original words from which the stem was derived such that the user is presented with a word and not a stem.

Single words are often not enough to convey the full meaning that the user wishes to explore. Most of the time, the user is looking for material that is characterized by a combination of two or more words rather than just one. For example, no one would contend that *Supreme Being* together is more semantically rich than *Supreme* and *Being* taken separately. The idea here is to establish the set of word combinations that reflect the meaning of the document.

In ESCAS, we thus extract collocations of 2 consecutive words. Collocations are of interest to us for two reasons. First, they are important for computational lexicography. Being a multi-word term combined together in a significant way, they are used to create a more useful index. Second, since the process of generating collocations uses a ranking scheme which integrates the importance of specific words within a document, this process produces a set of terms that characterize the documents. These terms can then be selected by the user to perform query rewriting (see section 3.3) with a single touch of a function key.

To identify collocations, we used the standard technique of Hypothesis testing with the Pearsons chi-square test described in [5]. Essentially, Pearsons chi-square test compares the observed frequency of pairs of words with the frequency expected if the two words were independent. For more technical details, the interested reader is referred to [5].

Figure 1 shows a screen shot of the system with a collocation-based index. The index term *age modern* is found twice in the text. By clicking on either excerpts *For two hundred years, people ...* or *They tackled all their practical and ...*, the user is taken directly in the right position in the document.

3.2 Paragraph Retrieval

As discussed in section 1.1, blind students have a difficult time skimming documents because they do not have an overall view of the text. They currently must read each document to determine their relevance, then select only a few and read them back again to analyze their content. The idea behind paragraph retrieval is to suggest to the user places where he would start reading that would be the

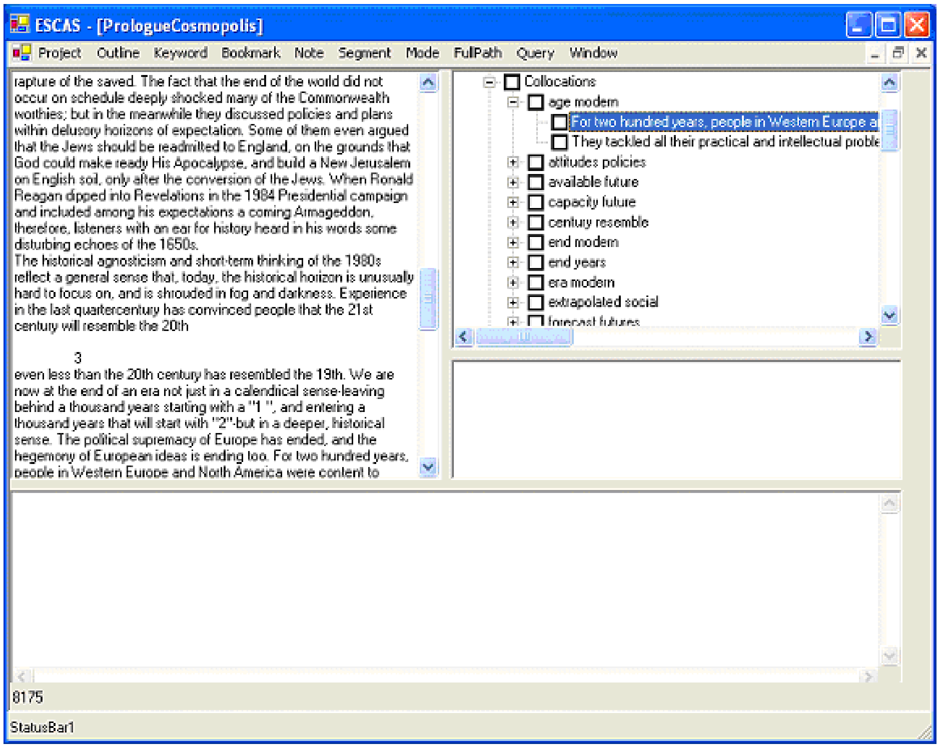

Fig. 1. Collocation-based index in ESCAS

most relevant to the subject being treated. If he could do a structured reading of the material, he would save time by having less to read or not having to read the material more than once. This problem can be reformulated as an Information Retrieval (IR) task where the query is the subject of the thesis or paper being written. We wish to provide a student with search facilities when reading a document and cater at the same time for the different levels of understanding of that student.

In ESCAS, the document being studied is divided into a set of paragraph vectors using the data accumulated during the parsing of the document for index terms (see previous section). We implemented the standard vector-space IR model with a tf-idf weighting scheme, where partial matching allows the ranking of the degree of similarity measured by the cosine of the angle between a paragraph vector and the query vector. Initially, the query is taken to be the topic of the composition. Only the paragraphs relevant to the topic (or those having the highest cosines) are retrieved.

3.3 Query rewriting and reformulation

The problem with IR is that often very pertinent paragraphs containing terms semantically close to the query (such as synonyms) are not retrieved while noise is introduced in the retrieval set with paragraphs that are not pertinent. As an example, let's assume that a Philosophy term paper has to be written on

Potentiality and Actuality: Aristotle's reintroduction of unrealized possibilities and reaffirmation of becoming. A query based on this subject would not include the word *Being* since that word is not among the words in the query. So the student has to somehow add this word to the query to improve the retrieval. Based on the new retrieval results, new words or phrases will suggest themselves (through the index for example) and the user will start exploring the material with different queries.

Figure 2 shows a screen shot of the system with paragraph retrieval. The benefit that can be derived from this feature is faster access to pertinent material within the text or exploring the material in unexpected but interesting directions. One must remember here that we are trying to find a substitute for skimming. Furthermore, this iterative reformulation exercise suggests that the problem can be viewed as having three dimensions or can be positioned within a three-dimensional space, the type of retrieval, the expansion of the query and the level of understanding of the user.

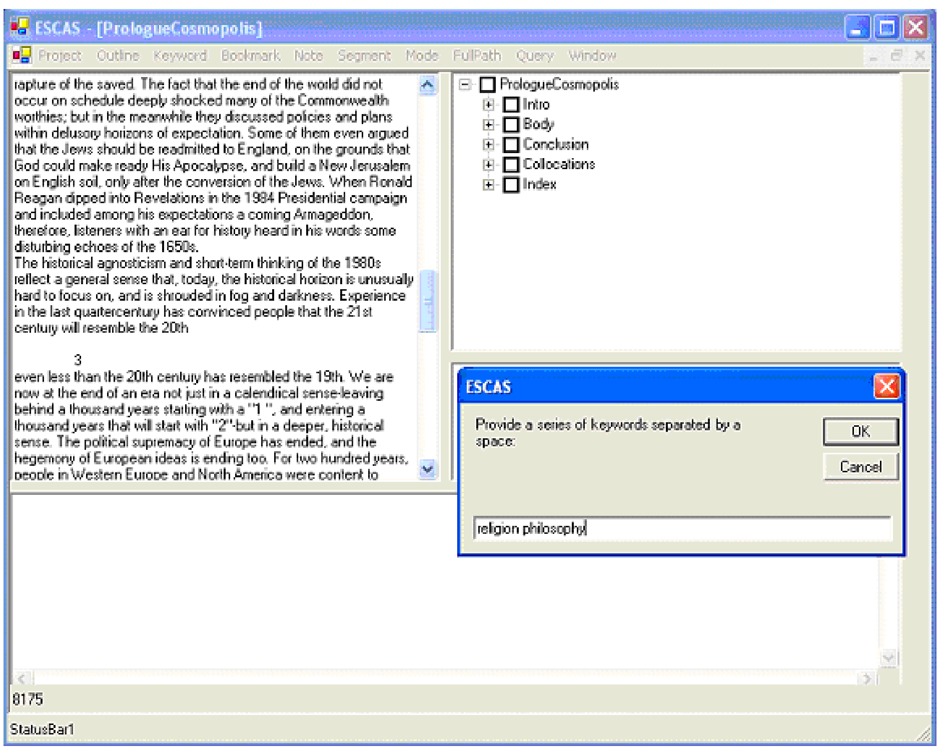

Fig. 2. Paragraph retrieval in ESCAS

WordNet Query Reformulation can quickly turn short if the user does not have some help in getting more ideas on the subject. WordNet was thus incorporated into the system. The idea was to make available to the user synonyms, hypernyms, hyponyms, and other lexically and semantically related terms. Their respective links allow the user to explore a certain domain and thus structure and enrich his query reformulation process. With a function key, ESCAS allows the user to expand his query from where the cursor is in the Wordnet data so he can resubmit that query within his exploration process. Figure 3 shows a small excerpt of the links available when a search is done on hyponyms of the word *Being* in WordNet.

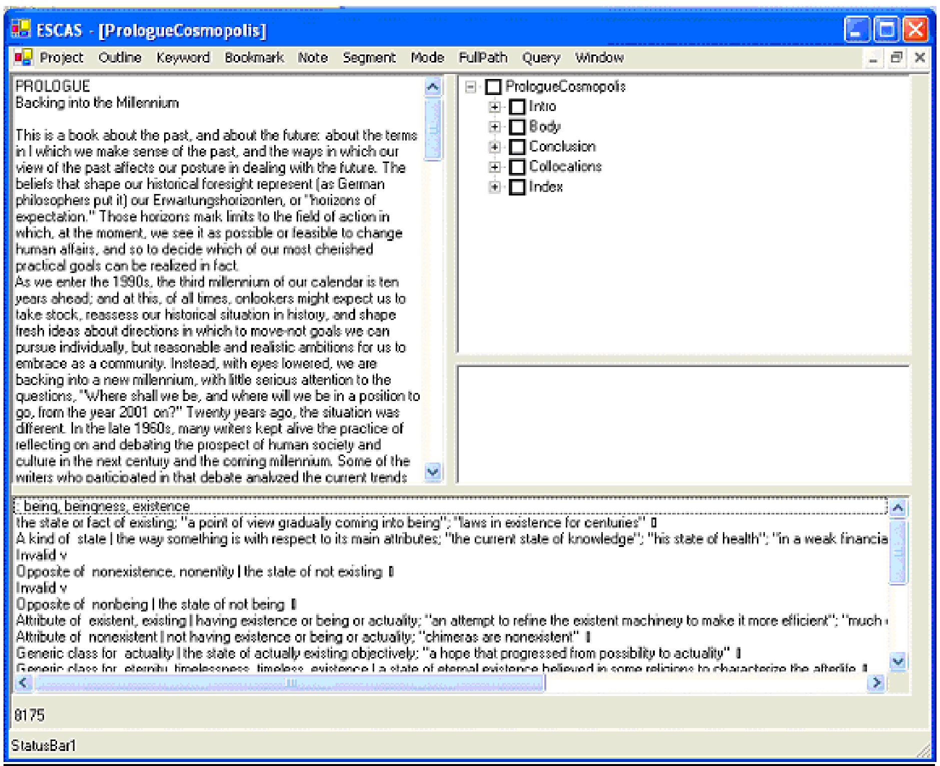

Fig. 3. Example of using WordNet in ESCAS

4 Evaluation

To evaluate the system, both a quantitative and a qualitative evaluation were performed, but with only one blind university student. We also showed the system to the coordinator of Concordia University's Office for Students with Disabilities - himself blind, and gathered his opinion. One of the major difficulties in performing the evaluation, is to find blind students who are willing and have the time to evaluate the system. Remember that blind students require more time to do their course work. Asking them to spend time on the evaluation of a prototype adds a load to their already busy schedule. Nevertheless, as discussed later, a more complete evaluation with several blind and non-blind post-secondary students is planned for.

4.1 Qualitative Evaluation

From the beginning, a blind student was involved in the project. Didier is now a 2^{nd} year Concordia University student enrolled in Philosophy and Religion. Since CECEP (Québec's pre-university level), he had to write compositions for his course work. He was involved in the project since 2001, helping us determine the requirements of blind students, and since 2002, Didier had been given various prototypes to try out and evaluate.

Through Didier's use of the system, we interviewed him and asked him to describe his experience. After some training period where he had some good and some bad experiences, Didier characterizes his experience as follows: *If you know where you are going, ESCAS is a fantastic and indispensable tool. If you don't, it is a bad tool. More concretely, you must first have a good idea of your topic and your material first before creating an outline and fleshing it out into a paper.* In the beginning he retrieved many documents and paragraphs, then spend a lot of time reviewing it all. He was falling into the familiar trap of beginning students who would highlight the whole book.

Today, Didier claims that it takes him about half the time to write a composition when using the system. He can skim literature more quickly and can concentrate on more creative thinking than he could without the system. All in all, Didier has been using the system every semester since we offered it to him and he considers it a necessary tool for reading and writing.

We also showed the system to the coordinator of Concordia University's Office for Students with Disabilities - himself blind. Although he did not use the system for an actual essay writing exercise, he was very enthusiastic about it and will recommend it to his peers at various local colleges (Dawson, Vanier, Marianopolis and Cégep du Vieux-Montréal).

4.2 Quantitative evaluation

From the beginning, the goal of the system was to reduce the time gap between blind and regular students in writing compositions. To measure the utility of the NLP tools, we therefore needed to compare the time it took a typical user to write a composition with and without the system. Dider feels that it now takes him about half the time to write an essay; but we wanted *hard* data to support this claim. The difficulty here, is that any person (hopefully students, too) get better at a task each time they perform it. We cannot ask the same person to write a composition on the same topic twice; the second time will surely be faster and better. We cannot compare different compositions written by the same person, as they may be of different levels of difficulty, different level of research will be involved ... Since Didier had kept record of all his course work since 2001, we then compared the time it took him to write compositions for his courses compared to the professors' requirements for regular students. In 2001, Didier did not have the software, in 2002 and 2003, he was given various prototypes that did not include NLP tools, and since 2004, he was given ESCAS. Table 1 shows the data for years 2001, 2004 and 2005.

Table 1. Time to write a composition with/without the system compared to professors' requirements

	Regular Students			**Didier's Actual Time**			**Time**	
Year	**Given**	**Due**	**Time**	**Given**	**Handed-in**	**Time**	**Ratio**	**Diff.**
2001 - w/o ESCAS	01-Mar	15-Apr	45 days	01-Mar	30-May	90 days	2.00	45
2001 - w/o ESCAS	09-Feb	30-Apr	81 days	09-Feb	27-Jun	138 days	1.70	57
2004 - w/ ESCAS	08-Sep	02-Dec	84 days	08-Sep	06-Dec	88 days	1.05	4
2004 - w/ ESCAS	08-Sep	13-Dec	95 days	08-Sep	19-Dec	101 days	1.06	6
2005 - w/ ESCAS	07-Feb	07-Mar	30 days	07-Feb	30-Mar	53 days	1.77	23
2005 - w/ ESCAS	16-Nov	09-Dec	23 days	02-Nov	12-Dec	40 days	1.74	17
2005 - w/ ESCAS	09-Nov	23-Nov	14 days	16-Nov	15-Dec	29 days	2.07	15

In 2001, it took Didier about twice as long to write a composition than regular students. If the professor gave 45 days to regular students, it would take him 90 days to achieve the same task. In 2004, the data seems to show a net reduction in time - he handed in his work about the same time as regular students. However, for 2005, it takes him twice as long again ... The data was inconclusive.

However, when looking at the difference between the time of a regular student and that of the evaluator, we see three clusters of data corresponding to each of the three years. ESCAS does show an improvement. The higher differences in 2005 compared to 2004, can be explained by more material having to be read and more courses taken simultaneously. Why would differences reflect reality better than ratios? Given the evaluator is adamant that he takes less time when using the application, we would like to suggest an explanation. The student doing the evaluation has a physical deficiency, not an intellectual one. Using a ratio would assume that every task takes more time affecting the overall proportion uniformly on all the variables. This would correspond more to an intellectual deficiency. In the case of a physical deficiency, accessing the material takes more time. But once the material is available, the creative and intellectual skills come into play. ESCAS does not help at all at the intellectual level. Considering the differences instead of the ratios amounts to looking at the extensions in time required to finish the paper. Let us apply this rule to the first and last entries. We can safely assume that the time to access the material for a sighted person is marginal when compared to the problems met by a blind person. Giving this access time a value of 1, we have 44 days and 13 days of intellectual work respectively. We also assume parity at the intellectual level. We have a resulting 46 days (90-44) and 16 days (29-13) of access time. This gives us a true ratio difference of 46.0 to 16.0 which is a huge difference when compared with the 2.0 and 2.07 pair.

To conclude, with the small amount of data analysed, we have not scientifically demonstrated that ESCAS helps. Future work definitely includes a formal evaluation of the system with several students, and compare their use of the system. The goal here is to make sure that we have not developed a system geared towards the personal preferences of one person, but that it is actually useful for most blind students. For this, we plan on asking the participation several blind and non-blind subjects from a post-secondary school. However, we doubt that a truly quantitative evaluation can be performed. As with any software, the perceived benefits may not correspond to the actual benefits. Overall, if the user freely chooses to use the system, it may be enough to declare it useful.

5 Conclusion and Future Work

In this paper, we presented ESCAS[5], a prototype system to assist the visually handicapped in writing compositions. The system can be seen as an information probing and gathering environment that offers features based on NLP techniques. The purpose here is to investigate where and how far existing NLP techniques can facilitate the access of students to information pertinent to their research. These techniques convey additional information about the content of a document and faster access to relevant positions within the document.

Currently, the evaluation of the system is more qualitative, than formal. Future work definitely includes a formal evaluation of the system with several students, and compare their use of the system. The goal here is to make sure that we have not developed a system geared towards the personal preferences of one person, but that it is actually useful for most blind students. For this, we plan on asking the participation several blind and non-blind subjects from a post-secondary school. However, we doubt that a truly quantitative evaluation can be performed. As with any software, the perceived benefits may not correspond to the actual benefits. Overall, if the user freely chooses to use the system, it may be enough to declare it useful.

Acknowledgment The authors would like to thank Didier and Leo for their time evaluating the system and for their constant feedback and encouragement. Many thanks also to the anonymous referees for their valuable comments on a previous version of this paper.

References

1. *CSUN's 19th Annual International Conference "Technology and Persons with Disabilities"*, Los Angeles, March 2004.

2. *CSUN's 20th Annual International Conference "Technology and Persons with Disabilities"*, Los Angeles, March 2005.
3. N. Cowan. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24:87–185, 2001.
4. Julia Hodges, Shiyun Yie, Ray Reighart, and Lois Boggess. An automated system that assits in the generation of document indexes. *Natural Language Engineering*, 2:137–160, 1996.
5. Christopher D. Manning and Hinrich Schütze. *Foundations of Statistical Natural Language Processing*. MIT Press, 1999.
6. G. A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63:81–97, 1956.
7. Fred Orelove and Lissa Power-deFur. *Inclusive Education: Practical Implementation of the Least Restrictive Environment*. Jones and Bartlett Publishers, 1997.
8. M.F. Porter. An algorithm for suffix stripping. *Program*, 14(3):130–137, 1980.
9. T. Shimizu, S. Smoliar, and J. Boreczky. AESOP: An Outline-Oriented Authoring System. In *Proceedings of the Thirty-first Annual Hawaii International Conference on System Sciences*, pages 207–215, Hawaii, January 1998.
10. N. Wacholder, D. K. Evans, and J. L. Klavans. Automatic identification and organization of index terms for interactive browsing. In *International Conference on Digital Libraries. Proceedings of the 1st ACM/IEEE-CS joint conference on Digital libraries*, pages 126–134, 2001.