

User Centered Evaluation of an Automatically Constructed Hyper-TextBook

FABIO CRESTANI

*Department of Computer and Information Sciences
University of Strathclyde
Glasgow, Scotland, UK
f.crestani@cis.strath.ac.uk*

SPYRIDON P. NTIOUDIS

*Department of Electrical and Computer Engineering
National Technical University of Athens
Athens, Greece
dioudis@softlab.ntua.gr*

As hypertext systems become widely available and their popularity increases, attention has turned to converting existing textual documents into hypertextual form. An important issue in this area is the fully automatic production of hypertext for learning, teaching, training, or self-referencing. Although many studies have addressed the problem of producing hyper-books, either manually or semi-automatically, the actual usability of hyper-books tools is still an area of ongoing research. This article presents an effort to investigate the effectiveness of a hyper-textbook for self-referencing produced in a fully automatic way. The hyper-textbook is produced using the Hyper-TextBook methodology. We developed a task-based evaluation scheme and performed a comparative user-centred evaluation between a hyper-textbook and a conventional, printed form of the same textbook. The results indicate that the hyper-textbook, in most cases, improves speed, accuracy, and user's satisfaction in comparison to the printed form of the textbook.

The motivation for the work reported in this paper can be found in the explosive use of hypertext systems for teaching, learning, training, or simply satisfying an information need. This technology is a relatively new arrival in educational settings, and evaluation studies conducted on hypertext systems' usability as educational tools produce contradictory results (Mayes, Kibby & Anderson, 1990). In addition, previous work in the evaluation of hyper-books has not provided researchers with conclusive evidence towards either superiority or inferiority of these books in comparison to conventional educational material. Even a brief review of this literature would require a much longer space than we have (Egan, Remde, Gomez, Landauer, Eberhardt, & Lochbaum 1989; Leventhal, Maynatt Teasley, Instone, Schertler Rohalman, & Farhart, 1993; Tebbutt, 1999). However, because of some promising evaluation results in favour of hypertext systems, the need for further evaluation arises.

The aim of the work reported in this article is to investigate the effectiveness of a particular model of a hyper-textbook directed specifically to the task of self-referencing. A *hyper-textbook* is the result of converting a textbook into its hypertextual form. A *textbook* is a book containing facts about a specific subject, and it is used by people studying that subject. Although a textbook is often used for learning and teaching, in many cases it is also used for reference and fact finding, in particular after the subject dealt by the textbook has already been studied. This is not only the case for teaching material, but also for technical manuals, documentation, and handbooks.

The hyper-textbook we used was developed using the Hyper-Textbook methodology described in Crestani & Melucci (1998a). This methodology enabled the production of a hypertext in a fully automatic way. The evaluation was conducted using a task-based methodology partially developed in the context of the Esprit Working group "Mira." We believe there is a great need for methodologies and tools for the fully automatic production of hypertext to be published on the Web. This is particularly true for documents that are used for self-referencing purposes, such as teaching material (notes, textbooks), software/hardware documentation, technical manuals, and so forth. Only an inexpensive production of such material will encourage its full availability on the Web.

The article is structured as follows. First, a brief introduction to the Hyper-TextBook system is provided. Then the experimental design adopted for the evaluation of a hyper-textbook is presented. Following the experimental design is the section devoted to the presentation and analysis of the experimental results. In the last section the conclusions of the work are reported and some issues around which future work is evolving is suggested.

THE HYPER-TEXTBOOK SYSTEM

As hypertext systems become widely available and their popularity increases, attention has turned to converting existing textual documents into hypertextual form. Many projects have focussed on such tasks, for example, SuperBook (Egan et al., 1989) and CORE (Egan, Lesk, Ketchum, Lochbaum, Remde, Littman, & Laundauer, 1991). Most of the research in this area has been concerned with the effective identification of inter-document and cross-document links to enable a user to access information in a nonlinear way (Frisse, 1988; Smeaton & Morrissey, 1995; Agosti, Crestani, & Melucci, 1997; Allan, 1997). The Hyper-TextBook project is not an exception.

The methodology developed in the context of the Hyper-TextBook project has a long history. In Agosti, Colotti, and Gradenigo (1991) a two-level conceptual model of hypertext called EXPLICIT was presented. This model enables the user to have a clear frame of reference during the browsing of a large structured hypertext built from a single large document or collection of documents. Later, Agosti and Crestani developed a methodology for the construction of a hypertext organised according to that model (1993). The methodology makes use of Information Retrieval (IR) techniques for extracting and linking semantic objects like index terms or parts of text. A few years later, TACHIR, a tool that automates that technique, was developed and tested (Agosti, Crestani, & Melucci, 1995).

In 1998, the Hyper-TextBook project in which a methodology and a tool for the automatic construction of hyper-textbooks from textbook (Crestani & Melucci, 1998a) were developed. The Hyper-TextBook methodology makes use of the experience gained with EXPLICIT and TACHIR to automatically produce a hyper-textbook written in HTML that can be browsed with any web browser. The construction of a hyper-textbook is carried out by the Hyper-TextBook tool in a batch way. Hyper-TextBook takes as input a number of files containing an electronic version of the textbook and of the subject index (accepting a number of different formats) and produces as output the entire hyper-textbook in a large number of HTML files. The entire process can be done off-line and requires only the specification of some parameters necessary for the evaluation of similarity between pages, between terms, and between terms and pages. Depending on the size of the textbook the automatic authoring of the hyper-textbook can require minutes or hours, but does not require any human intervention.

A full description of the project is outside the scope of this article. Here we will briefly present an example of the use of a hyper-textbook produced using the Hyper-TextBook methodology. The example describes how a user

can browse and search a hyper-textbook, the same used for the evaluation reported in this article. This hyper-textbook, *Information Retrieval* by Van Rijsbergen (1979), is available at <http://www.dei.unipd.it/~ims/htb/>.

Example

Figure 1.a shows the home page of the hyper-textbook. Three entry points for browsing are available. The first one allows access to the hyper-textual version of the subject index, from which the user can go to a term. The second entry point corresponds to the list of pages. The third entry point permits access to the table of contents. We assume here that a user goes directly to the subject index where he/she selects a particular term that best represents what he/she is looking for.

Figure 1.b is the page for the term “cluster profile” of the subject index. From this page the user can access:

- the list of pages relevant to the current term (according to its patterns of occurrences in the pages);
- the list of terms similar to the current term (according to the pattern of co-occurrence of terms in the textbook); and
- other terms linked through the manually inserted “see also” links (according to the author’s subject index).

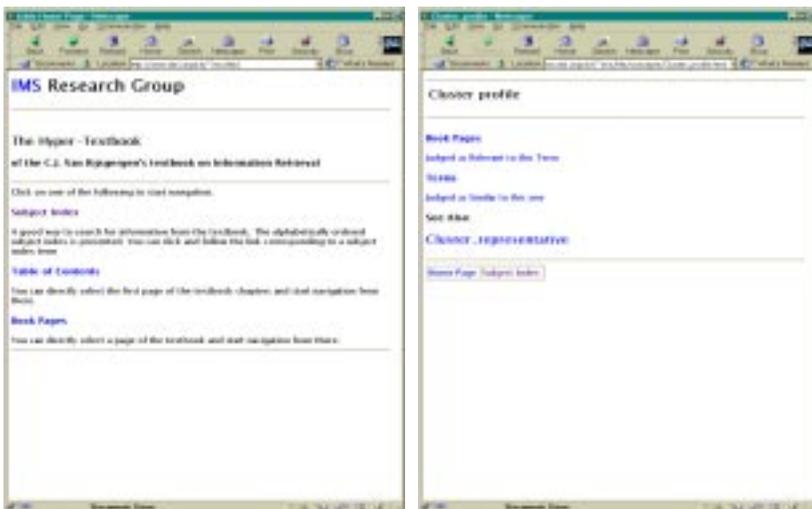


Figure 1. a) Home page of the hyper-textbook. b) Page of a term of the subject index.

Suppose the user chooses to see the list of pages relevant to the term. Figure 2.a reports the list of textbook pages ranked relevant to the term “cluster profile.” As this term was not associated to any pages by the author, the list of linked pages was built using the pages linked to the term “cluster representative.” It is worth noticing that the top ranked page is one of those that was pointed out as relevant by the textbook author in the subject index, but the following one, for example, was not. This means that the technique used to infer semantic links is able to preserve the author’s judgements of relevance, as well as to identify some additional relevant pages (Crestani & Melucci, 1998b). Links are ranked according to the normalised term weight within the textbook page.

Figure 2.b contains page 51 of the textbook. The user arrived at that page just by clicking on the page number. The hyper-textbook displays, for each page, the page number, a summary, and the full text of the page. Both at the top and at the bottom of each page some buttons are located that link the current page to the next page, to the previous page, to terms describing the content of the page, and to similar pages. Links to the next and previous pages enable a linear reading of the textbook. Links to terms describing the content of the page or to similar pages enables a non-linear navigation of the hyper-textbook according to the EXPLICIT model.

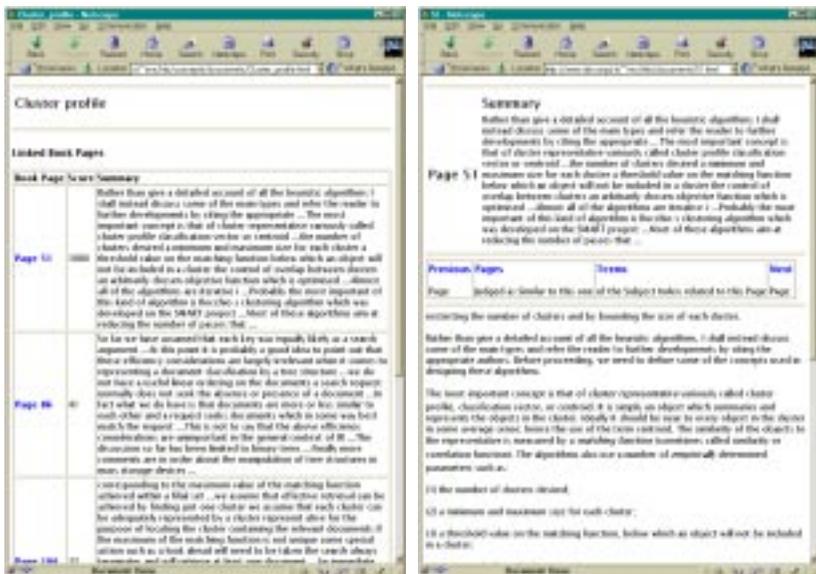


Figure 2. a) List of textbook pages relevant to the term “cluster profile.” b) A page of the textbook relevant to the term “cluster profile.”

Figure 3.a shows the list of similar pages linked to the current page. Links are ranked according to a normalised page similarity measure based on the distribution of terms in the pages.

The user can also see the terms that are similar to a selected term, as depicted in Figure 3.b, which reports the list of terms similar to the term “cluster profile.” Links are ranked according to a normalised term similarity measure.

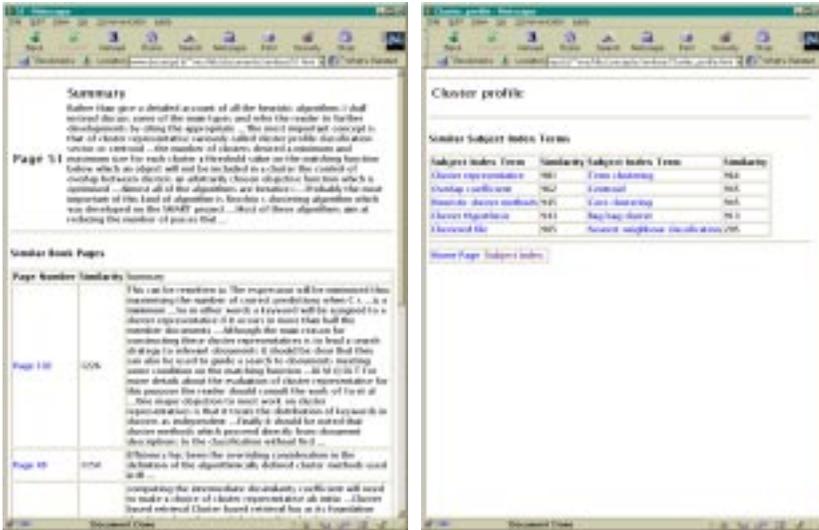


Figure 3. a) List of pages similar to page 51. b) List of terms similar to the term “cluster profile.”

EVALUATION DESIGN

Evaluation is a very important issue in the area of hypertext construction and hypertext usability. This motivated us to carry out an in-depth evaluation of the Hyper-TextBook system as a fully automatic way of electronically authoring hyper-textbooks from textbooks. At this stage, we decided to concentrate our evaluation on the usability of the final product, the hyper-textbook, instead of the efficiency of the automatic authoring process.

Research Hypothesis

We made the following experimental hypothesis: *The hyper-textbook is a more effective tool for seeking information than the original printed form of the textbook.*

The textbook chosen for the testing of the above hypothesis was Van Rijsbergen's *Information Retrieval* (2nd edition), henceforth referred to as the textbook. The primary reason behind this choice can be ascribed to the fact that Van Rijsbergen, who owns the copyright of the book, granted us permission for its use for the evaluation reported in this article.

Evaluation Scenario

Once the research hypothesis has been established the discussion should turn to the basic design settings upon which the testing of the hypothesis will be carried out. In this case, the independent variable has two levels: the presence (hyper-textbook) or absence (printed textbook) of hypertext functionalities. Participants in the evaluation have to perform some tasks with and without the assistance of hypertext links and of hyper-textbook additional features. The dependent variable is the effectiveness of the participants in completing these tasks. The primary objective is to prove that any variation of the effectiveness between participants is attributed only to the change in the level of the independent variable, that is, to the additional functionalities enabled by the hyper-textbook.

Since the independent variable has two levels, the same number of groups of participants will be employed. One group will perform some tasks using the hyper-textbook, while the other group will perform the same tasks using the printed form of the textbook. In our case, due to restrictions on both time and resources, each group consisted of 10 participants, arguably, a number deemed adequate to produce significant experimental results. The participants were postgraduate students doing a conversion course in Computer Science with little, or no, experience in the subject of the textbook.

The issue of keeping the irrelevant variables under control while testing the validity of a research hypothesis is crucial. For the control of irrelevant variables the independent groups design was adopted.

The Experiment

Each participant was assigned to one of the two levels of the independent variable and was presented with seven questions. The same questions were given to all participants regardless of the group they belonged to. Questions were of two categories: open-ended and close-ended. The close-ended ones were supplemented by a list of four answers of which only one was correct. Questions included terms occurring in the subject index, and/or in the table of contents and or in paragraph titles. The answers to the questions used in the experiments were given by the author of the textbook. This evaluation procedure is supposed to mimic a student using the book for some homework or examination. For more details about the experiment and the actual questions see Ntioudis (1998).

Unknown to the participants, some of the questions were designed to be well suited for answering using the conventional, printed textbook, while the rest exploited, partially or fully, the functionality of the hyper-textbook. The two groups of questions were of equal size.

After a short practice session, each participant had a time limit to answer each of the seven questions. Participants were allowed to type in "I don't know" as an answer, at any time before the end of the specific time limit, if they felt that they could not find the answer. Apart from the participants' time for completing the assigned task, the experimenter took note of the steps taken by the participants to accomplish the task of answering each question. This was done to reveal participants' strategies in completing the assigned task. At the end, each participant was asked to complete a questionnaire, and after that, a brief discussion took place to gather comments on the experimental procedure.

EVALUATION

As stated previously, the dependent variable the authors wished to examine through experimentation was the effectiveness of the users in the process of satisfying an information need by searching in a textbook. To achieve this goal, a set of criteria covering the aspects of the dependent variable had to be defined. These criteria were:

- the accuracy of the answers;
- the speed with which the search for the answers was completed; and
- the subjective opinion of the users about the assistance provided by the different forms of the textbook in completing the task assigned to them.

Apart from establishing the research hypothesis, the purpose of the set of criteria was to draw conclusions about the effectiveness of the hyper-textbook in the task of aiding a user satisfying an information need. Since the evaluation of the hyper-textbook is substantially user-centred, the performance of the users under such task-based environment can be an indication of the effectiveness of the hyper-textbook.

Accuracy

The actual performance of the users in the task of satisfying an information need must be measured. In this case, an obvious way to do so was to take into account the “correct answers” found by the participants. However, since the questions were of two different types (open-ended and close-ended), issues of judging their correctness arise.

Assessing the correctness of the answers for close-ended questions is quite easy. On the other hand, assessing the correctness of open-ended questions is more complicated. Since open-ended questions do not have a predefined answer, the participant may give the answer which, according to his/her opinion, thinks is the correct one. Answering open-ended questions involves composition and is a more creative process than just selecting an answer from a list of predefined answers. Therefore, the task of assessing the correctness of such questions requires adopting a scoring method to apply to each answer. A list of terms relevant to each open-ended question was developed. Copies of the answers were then given to an assessor who scored them “blindly” with the lists of terms. In this case, the success rate for each question was defined as the ratio (in %) of the number of terms contained both in the participant’s answer and in the list. In any case, no answer to a question or “I don’t know” was considered as a wrong answer and the success rate in answering that question was set to 0%.

Speed

The experimenter wrote down the time the participant spent in answering each question. This measure was used as an indication of the speed with which the participants performed the assigned tasks. It provided evidence about which experimental condition facilitated the faster execution of the user-performed task, although other factors may affect that condition, like, for example, the familiarisation of each participant with the interface and functionalities of the hyper-textbook.

Subjective Opinion of the Participants

In both experimental groups, participants were asked to state their opinion about the difficulty in answering the questions. In the group using the hyper-textbook, participants were also asked to state their opinion about the difficulty of navigating in the hyper-textbook. The information was collected from the questionnaire that each participant completed. In both cases, difficulty was measured on a scale from 1 (Very easy) to 5 (Extremely difficult). The questionnaire for the group using the hyper-textbook also contained a question referring to what the users found most difficult in using the hyper-textbook. Finally, two questions on whether the time to answer the questions was enough and on what would help participants answer questions in a more efficient way, were in every questionnaire. Although the opinions of the participants, as expressed in the final three questions and in the postexperimental discussion sessions, cannot be quantified, they can provide interesting and useful insights into the utilisation of the two different forms of the textbook.

RESULTS

For convenience, the group of participants assigned to work with the hyper-textbook will be referred to as group H, while the group of participants that performed the experiment using the conventional, printed textbook will be known as group P.

ACCURACY

The data presented in Table 1 have been acquired by averaging the results of the success rate of the answer to each question over the total number of answers, thus producing the average success rate values per question type. The success rate for group H in answering the close-ended questions was marginally greater than that for group P. On the other hand, for the open-ended questions, the success rate of group P was a bit higher than the success rate of group H. To establish the statistical significance of these results, *t*-tests were performed, indicating that, with probability of error 0.05, the results are attributed to the change of level of the independent variable and not to chance factors.

Table 1
Success Rate Values (in%) for Answering Questions

<i>Type of questions</i>	<i>Group H</i>	<i>Group P</i>
Close-ended questions	95	90
Open-ended questions	63.2	68.4

All participants from group H provided an answer to the open-ended questions while participants from group P failed to do so in two cases. By “failed to provide an answer” it is implied that no answer to a specific question was given or the participant typed in “I don’t know.” This means that, for the open-ended questions, group H answered (either correctly or not) 4% more questions than group P. Nevertheless, group P had a higher success rate than group H which means that answers from group H were not as complete and correct as answers from group P.

Speed

Table 2 presents the average speed results. For close-ended questions, group P was marginally faster than group H, although this difference is not statistically significant. For open-ended questions, however, results were quite the opposite. The fastest group was group H with an average time of 279.9 sec. for answering each question, while group P had a average time of 311.9 sec. *T*-tests attributed these results to the change of level of the independent variable, with probability of error 0.05.

Table 2
Speed Results (in seconds) for Answering Questions

<i>Type of questions</i>	<i>Group H</i>	<i>Group P</i>
Close-ended questions	113.4	108.7
Open-ended questions	279.9	311.9

Subjective Opinion of the Participants

Regarding the perceived difficulty of the questions, Table 4 shows that 82% of the participants in group H think that questions were relatively easy to answer. The same opinion was expressed by 90% of the participants in group P.

As far as time to answer is concerned, roughly 90% of the participants in both groups thought the time was enough (Table 3). Bearing in mind that participants were nonexperts in the field of Information Retrieval, the results are indicative of the aid of the medium in answering the questions.

Table 3
Participants Opinion on Whether There was Enough Time to Answer the Questions

<i>Was the time enough?</i>	<i>Group H</i>	<i>Group P</i>
Yes	90	89
No	10	11

Table 4
Participant Opinion (in %) About the Difficulty of the Questions

<i>Was the time enough?</i>	<i>Group H</i>	<i>Group P</i>
1 (very easy)	10	9
2	2	0
3	70	80
4	18	11
5 (very difficult)	0	0

In addition, participants in group H were asked to rate navigation in the hyper-textbook on a scale from 1 (Very easy) to 5 (Extremely Difficult) when completing their questionnaires. The results depicted in Table 4 indicated that only 18% found it difficult. This means that most of the participants clearly understood EXPLICIT and the navigation of the hyper-textbook.

Participant Strategies

Participants' strategies in answering the questions using the hyper-textbook comprised an extensive use of, primarily, the Subject Index and, secondary, the Table of Contents. In fact, 80% of the participants in group H started their search from the subject index. Every time a specific term from the subject index was selected, the next step was choosing the link that led to pages judged as similar to that term. Most of the time (70%) the link that lead to terms judged as similar to a previously selected term from the Subject Index, was not used.

By selecting the link leading to similar pages, a list of similar pages was presented to the user. Only 10% of the participants examined all the pages in that screen. The first three pages with the higher scores were the ones examined by the majority of users. Participants using the printed form of the textbook adopted an obvious strategy: look for a specific term in the subject index unless the question explicitly identified the chapter where the answer could be found.

CONCLUSIONS AND FUTURE WORK

As expected, the results of our experimentation did not provide definitive indications over the refutation or approval of the research hypothesis under testing. The outcome of the analysis of the experimental data was that only in some of the experimental tasks the hyper-textbook was superior in matters of accuracy and speed. In a number of tasks, participants using the printed textbook performed better than participants using the hyper-textbook. The advantage of each form of the textbook over the other was, in most cases, marginal and in some cases not statistically significant.

In addition, the interpretation of the differences in effectiveness of some experimental conditions is problematic, to say the least. Why hyper-textbook users were slightly better in speed for open-ended questions and slightly better in accuracy for close-ended ones? What else can we gather from the subjective impressions of users, a part from the obvious? More evaluation work is definitely needed.

Nevertheless, despite its clear limitations, the evaluation enabled us to identify some ways of improving the usability and the effectiveness of the hyper-textbook. These directions pertain to both the re-design of the hyper-textbook and its re-evaluation, following a formative design evaluation approach. In particular:

- the need for a better interface, reflecting in a better way the underlying model of the hyper-textbook (which was not understood by a small number of users), but maintaining the appearance of the printed textbook (so that users can locate quickly both the Subject Index and the Table of Content);
- better formatting of the text in each page to improve readability (more similar to a page of a book than to a Web page); particularly important was considered the highlighting of terms when a text page was reached from a term page;

- the need for the integration of a global search function in the hyper-textbook, to enable users to locate a section in the book pertinent to an information need without browsing (as requested by a number of IT minded users); and
- the presence of unnecessary and confusing features that needed to be removed, for example, the page summary.

A new version of Hyper-Textbook taking these results into account has been produced and is available at <http://www.cs.strath.ac.uk/~fabioc/htb/>. Even though this version is currently being used successfully as a teaching tool, a new evaluation is currently under way to reassess the usability of the constructed hyper-textbooks and continue to improve it.

Another important line of research that we are pursuing and that originated from the results of this study is related to the integration of “appearance” and “functionalities” in the authoring of hyper-textbooks. In the Visual Book project (Landoni, 1997), Landoni explored the importance of the visual component of the book metaphor in the production of “good” electronic books. The original aspect of the Visual Book project was the importance given to the visual components of the physical book when designing electronic books together with the interpretation of an electronic book as part of an electronic library intended as an informative system with specific and innovative features. The Visual Book was studied within the context of an electronic library by following an original approach which has highlighted and exploited its relation with the real object it imitates. In particular a new aspect, which has not been exploited to date, is that of *visual rhetoric* which is important to the design of both paper and electronic books.

The definition of visual rhetoric is tied, at least in a first instance, to the concept of text rhetoric, from which it is derived, as this is a well-established and popular example of rhetoric applied to written information. The idea is to define those parts that are more important for the comprehension of the meaning of the text. To achieve this purpose verbal and/or graphical techniques can be used. In fact, a document can be interpreted as a visible representation of a text according to its semantic contents (Southall, 1989). Thus visual rhetoric is simply the translation into graphical terms of the text rhetoric which results from both the logical structure of the text and its pragmatic component. It provides the reader with a graphical mark up language that is immediately recognisable on the basis of previous reading activity. Different graphical presentations suggest different readings and affect deeply the interpretation of the contents of the same text. These observations lead one to conclude that visual rhetoric is a crucial aspect for both

reading and browsing a document, as the findings of the Visual Book project have confirmed.

From the previous brief description of the Visual Book project, it is clear that the Visual Book and the Hyper-TextBook have a lot of similarities when it comes to their approach to electronic publishing. They both take the readers as the centre of the whole system and pay particular attention to their real needs. Both systems are the result of a careful study on what sort of electronic publications can be of real use to a specific category of users (scientists for the Visual Book, students for the Hyper-TextBook) and what additional features they may require or take profit from in the electronic version. A difference between the two projects is that while the Visual Book concentrates on the presentation issues, by looking for new presentation paradigms suitable for the electronic media in the attempt to improve the visual quality of the electronic version to be at least comparable with the paper one, the Hyper-TextBook focuses instead on providing additional features related to the electronic media in order to make the electronic version more flexible and powerful than the paper one.

It looks reasonable to expect that the Visual Book and the Hyper-TextBook approaches together could result in an enhanced *Visual Hyper-TextBook*. This could be achieved by providing the Hyper-TextBook with a visual interface that will take advantage of the Visual Book experience and its positive findings. The Visual Hyper-TextBook will then offer a very rich environment to study further issues related to the production of electronic books. We are already working in this direction (Landoni, Crestani, & Melucci, 2000).

References

- Agosti, M., Colotti, R., & Gradenigo, G. (1991). A two-level hypertext retrieval model for legal data. In *Proceedings of ACM SIGIR*, pp. 316-325, Chicago, IL, USA.
- Agosti, M., & Crestani, F. (1993). A methodology for the automatic construction of a hypertext for information retrieval. In *Proceedings of the ACM Symposium on Applied Computing*, pp. 745-753, Indianapolis, IN, USA.
- Agosti, M., Crestani, F., & Melucci, M. (1995). Automatic authoring and construction of hypertext for Information Retrieval. *ACM Multimedia Systems*, 3(1), 15-24.
- Agosti, M., Crestani, F., & Melucci, M. (1997). On the use of information retrieval techniques for the automatic construction of hypertexts. *Information Processing and Management*, 33(2), 133-144.

- Allan, J. (1997). Building hypertext using information retrieval. *Information Processing and Management*, 33(2), 145-159.
- Crestani, F., & Melucci, M. (1998a). A case study of automatic authoring: from a textbook to a hyper-textbook. *Data and Knowledge Engineering*, 27(1), 1-30.
- Crestani, F., & Melucci, M. (1998b). A methodology for the enhancement of a hypertext version of a textbook by the automatic insertion of links in the subject index. In *Proceedings of the IEEE ADL '98 Conference*, (pp. 157-166), Santa Barbara, CA, USA.
- Egan, D., Lesk, M., Ketchum, R., Lochbaum, C., Remde, J., Littman, M., & Landauer, T. (1991). Hypertext for the electronic library? CORE sample results. In *Proceedings of ACM Hypertext '91*, (pp. 299-312), San Antonio, TX, USA.
- Egan, D., Remde, J., Gomez, L., Landauer, T., Eberhardt, J., & Lochbaum, C. (1989). Formative design-evaluation of SuperBook. *ACM Transactions on Information Systems*, 7(1), 30-57.
- Frisse, M. (1988). Searching for information in a medical handbook. *Communications of the ACM*, 31(7), 880-886.
- Landoni, M. (1997). *The visual book system: a study of the use of the visual rhetoric in the design of electronic books*. Unpublished doctoral dissertation, Department of Information Science, University of Strathclyde, Glasgow, Scotland, UK.
- Landoni, M., Crestani, F., & Melucci, M. (2000). The visual book and the hyper-textbook: Two electronic books one lesson? In *Proceedings of the RIAO Conference*, (pp. 247-265).
- Leventhal, L., Mynatt Teasley, M., Instone, K., Schertler Roland, D., & Farhat, J. (1993). Sleuthing in HyperHolmes: An evaluation of using hypertext vs. a book to answer questions. *Behaviour & Information Technology*, 12(3), 149-164.
- Mayes, T., Kibby, M., & Anderson, T. (1990). Learning about learning from hypertext. In Jonassen, D. & Mandl, H. (Eds.), *Designing hypermedia for learning*, NATO ASI Series, (pp. 227-250). Heidelberg, Germany: Springer-Verlag.
- Ntioudis, S. (1998). *User centred evaluation of an automatically constructed hyper-textbook*. Unpublished M.Sc. dissertation, Department of Computing Science, University of Glasgow, Glasgow, Scotland.
- Smeaton, A., & Morrissey, P. (1995). *Experiments on the automatic construction of hypertext from text*. Technical Report CA-0295, School of Computer Application, Dublin, Ireland.
- Southall, R. (1989). Interfaces between the designer and the document. In J. Andre, R. Furuta, & V. Quint, (Eds.). *Structured documents*, (pp. 119-131). Cambridge, UK: Cambridge University Press.

- Tebbutt, J. (1999). User evaluation of automatically generated semantic hypertext links in a heavily used procedural manual. *Information Processing and Management*, 35(1), 1-18.
- Van Rijsbergen, C.J. (1979). *Information retrieval (2nd ed)*. London, UK: Butterworths.

Acknowledgements

The authors would like to thank a number of people: Keith van Rijsbergen for making available the textbook used in this experimentation, Massimo Melucci, Silvia Gabrielli, Monica Landoni, and Jane Reid for their useful suggestions on the evaluation methodology.

Spyridon Ntioudis carried out this work as partial fulfilment of the requirements for a M.Sc. in Advanced Information Systems at the University of Glasgow under the supervision of Fabio Crestani. The design and implementation of Hyper-TextBook were carried out by Fabio Crestani and Massimo Melucci at the University of Padova, Italy.