

Help Me Read! Expanding Students' Reading with Wikipedia Articles*

Arun-Balajee Lekshmi-Narayanan, Khushboo Thaker, Peter Brusilovsky, Jordan Barria-Pineda
School of Computing and Information
135 N Bellefield Avenue
Pittsburgh, PA, USA
{arl122,kmt81,peterb,jab464}@pitt.edu

ABSTRACT

In this demo paper, we present an implementation of an intelligent digital textbook integrated with external readings for students, such as Wikipedia articles. Our system applies the ideas of concept extraction from a digital textbook on topics in cognitive psychology and computer science for a graduate class in a large US-based university to generate search terms that can link with Wikipedia articles. Finally, we integrate these articles into the textbook reading interface, enabling students to quickly refer to Wikipedia articles in connection with the reading material of the course to understand a concept or topic that they struggle with or are interested in exploring further. With this demo, we present a system that can be utilized for data collection in a real-world classroom setup.

Keywords

Intelligent Textbooks, Digital Reading Systems, Wikipedia, Concept Extraction, Data Collection

1. INTRODUCTION

The rapid development of science and technology created a problem for college instructors who want to ensure that students receive up-to-date knowledge of the subject. While in the past, textbooks served as a predominant source of class readings, they frequently lagged behind the state-of-the-art. At present, many courses, especially at the graduate level, use a collection of recent research papers rather than textbooks as course readings. Unlike textbooks, which introduce domain knowledge gradually, taking care to explain critical concepts, research papers are written for audiences who are already familiar with core domain knowledge. Hence, research papers are challenging to read for unprepared students. Several authors have suggested that recommending relevant Wikipedia articles to explain complicated concepts

*(Does NOT produce the permission block, copyright information nor page numbering). For use with ACM_PROC_ARTICLE-SP.CLS. Supported by ACM.

A.-B. Lekshmi-Narayanan, K. Thaker, P. Brusilovsky, and J. Barria-Pineda. Help me read! expanding students' reading with wikipedia articles. In M. Feng, T. Käser, and P. Talukdar, editors, *Proceedings of the 16th International Conference on Educational Data Mining*, pages 525–528, Bengaluru, India, July 2023. International Educational Data Mining Society.

© 2023 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.
<https://doi.org/10.5281/zenodo.8115760>

could facilitate reading [1, 4]. Moreover, as an added benefit, the recommendations could make reading more personalized by encouraging students to explore readings related to their interests. However, implementing Wikipedia recommendations is not straightforward, since only some of the “concepts” mentioned in a research paper are useful recommendations in the context of a specific course. In this demo, we present a course reading system for research papers that uses advances in text mining to recommend the most relevant Wikipedia pages for every page of assigned readings. The system was tested in a full-term graduate course, where we also collected student feedback on the relevance and difficulty of recommended Wikipedia articles.

2. A READING SYSTEM WITH WIKIPEDIA RECOMMENDATIONS

To explore the opportunity to extend online reading with Wikipedia articles, we modified an online digital textbook reading platform, ReadingMirror [2], customizing it to research paper readings. The modified system inherited several useful features from the digital textbook platform, such as a table of contents (now *course reading plan*), annotations, and social comparison (Fig. 1). To extend the reading system with the recommendations of Wikipedia articles, we used text mining to extract entities from each reading page (see Section 3). Page-level extraction was used to provide recommendations on the page where the relevant concept is mentioned. Recommendations are provided using an expandable tab on a page margin. Clicking on this tab reveals a list of links to recommended articles which could be opened next to the article page. For example, if a page of an assigned article mentions “Allen Newell”, it is recognized as a useful Wikipedia concept and a link to the Wikipedia article is offered on Allen Newell, along with other recommendations for further exploration and reading (Fig. 2).

To instrument the classroom study reviewed below, all student work with recommendations (opening, scrolling, and closing the recommendation tab) is logged. In addition, we provide a simple interface for students to rate videos on the relevance and difficulty of recommended Wikipedia articles (bottom left in Fig. 2). To encourage ratings, the list of Wikipedia articles that the student has rated or read appears in a separate tab above the Wikipedia links tab.

3. ENTITY EXTRACTION

Previous work on Wikipedia linking compared the content of the page in the textbook that the student reads with the rel-

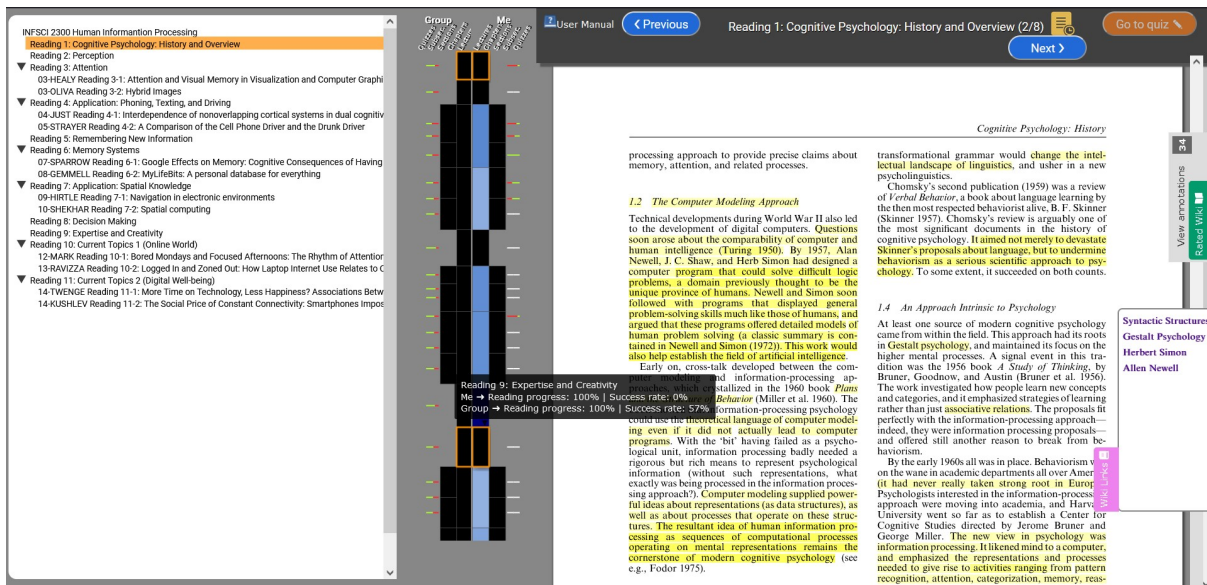


Figure 1: The interface of the reading system, Reading Mirror, with the course reading plan on the left and a page of the assigned reading on the right. A tab on the right of the reading page shows a list of recommended Wikipedia articles related to this page.

evant Wikipedia articles [1, 4]. However, these approaches could be noisy and generate relatively few recommendations. Since one of the goals of our project was to explore the feasibility of generating personalized recommendations that could engage students with different interests, we attempted to generate a somewhat excessive number of recommendations targeting the most relevant concepts mentioned on each page. To achieve this goal, we combined automatic concept extraction with heuristic filtering and embedding-based ranking for each reading page.

The first step in this process is to find Wikipedia concepts and entities mentioned on the target page. For each reading page, we extracted the entities mentioned on the page using the DBpedia Spotlight API¹. DBpedia Spotlight generates a list of entities in the submitted text along with corresponding Wikipedia pages linked to those entities. This list is usually large and noisy, so it requires post-processing. In the first step of post-processing, we filtered this list based on the semantic types of these entities, removing several irrelevant types of entities such as 'Event', 'Website', 'Film', 'Location', and 'Country'. We also removed entities that did not have a corresponding Wikipedia page in English. After the cleaning, we ranked the remaining entities. Since DBpedia Spotlight does not rank entities according to their relevance to the target page, we used the EMBED Rank [3]. For ranking with EMBED rank, we generated embeddings of the text on the page for which the recommendation is generated and the first paragraph in the ranked Wikipedia page. Top-*N* Wikipedia pages were recommended to the students.

4. A CLASSROOM DEPLOYMENT

To assess the usefulness of our idea and the quality of generated recommendations, we deployed the system as the course

processing approach to provide precise claims about memory, attention, and related processes.

1.2 The Computer Modeling Approach

Technical developments during World War II also led to the development of digital computers. Questions soon arose about the comparability of computer and human intelligence (Turing 1950). By 1957, Alan Newell, J. C. Shaw, and Herb Simon had designed a computer program that could solve difficult logic problems, a domain previously thought to be the unique province of humans. Newell and Simon soon followed with programs that displayed general problem-solving skills much like those of humans, and argued that these programs offered detailed models of human problem solving (a classic summary is contained in Newell and Simon (1972)). This work would also help establish the field of artificial intelligence.

Early on, cross-talk developed between the computational and information-processing approaches. The seminal work in the 1960 book *Plans and the Structure of Behavior* (Miller et al. 1960). The information-processing psychology approach to human problem solving was not actually lead to computer programs. With the 'bit' having failed as a psychological unit, information processing badly needed a rigorous but rich means to represent psychological information (without such representations, what exactly was being processed in the information processing approach?). Computer modeling supplied powerful ideas about representations (as data structures), as well as about processes that operate on these structures. The resultant idea of human information processing as sequences of computational processes operating on mental representations remains the cornerstone of modern cognitive psychology (see e.g., Fodor 1975).

Cognitive Psychology: History

transformational grammar would change the intellectual landscape of linguistics, and usher in a new psycholinguistics. Chomsky's second publication (1959) was a review of *Verbal Behavior*, a book about language learning by the then most respected behaviorist alive, B. F. Skinner (Skinner 1957). Chomsky's review is arguably one of the most significant documents in the history of cognitive psychology. It aimed not merely to devastate Skinner's proposals about language, but to undermine behaviorism as a serious scientific approach to psychology. To some extent, it succeeded on both counts.

1.4 An Approach Intrinsic to Psychology

At least one source of modern cognitive psychology came from within the field. This approach had its roots in Gestalt psychology, and maintained its focus on the higher mental processes. A signal event in this tradition was the 1956 book *A Study of Thinking*, by Bruner, Goodnow, and Austin (Bruner et al. 1956). The work investigated how people learn new concepts and categories, and it emphasized strategies of learning rather than just associative relations. The proposals fit perfectly with the information-processing approach—indeed, they were information processing proposals—and offered still another reason to break from behaviorism.

By the early 1960s all was in place. Behaviorism on the wane in academic departments all over America (it had never really taken strong root in Europe). Psychologists interested in the information-process approach were moving into academia, and Harvard University went so far as to establish a Center for Cognitive Studies directed by Jerome Bruner and George Miller. The new view in psychology was information processing. It likened mind to a computer, and emphasized the representations and processes needed to give rise to activities ranging from pattern recognition, attention, categorization, memory, reas-

¹<https://github.com/dbpedia/spotlight-docker>

reading system in a graduate course on human information processing in a large US-based university. In this lecture-based course, students were requested to read one or two assigned research articles prior to each lecture to prepare for a discussion. In the earlier offerings of this course, the articles were distributed to students in PDF form through a learning management system. In our study, the same articles were provided to students through the course reading system, which allowed us to generate a large number of page-level Wikipedia article recommendations for each assigned research article. The class had 11 lectures with a total of 17 research articles assigned for the required readings. The pages of these articles provided recommendations for 1,238 concepts linked to Wikipedia articles. As part of the learning process, we asked students to read at least 3 Wikipedia articles each week, selecting the most interesting ones for them from the set of recommended articles. In turn, to select these three most interesting articles, students were instructed to examine and rate (by relevance and difficulty) at least 10 recommended articles each week. For this work, students could earn up to one course credit point.

5. PRELIMINARY RESULTS

We collected learning data from 42 students enrolled in the class. In total, 772 out of 1238 recommended concepts linked to Wikipedia articles were explored and rated by students. An average of 12 students (*mean* = 12.73, *std* = 8.73) rated each concept for difficulty and 13 students (*mean* = 13.05, *std* = 9.05) for relevance. The 10 most popular concepts rated for relevance and those rated the most difficult are shown in Table 1. Since the students were guided by their interests, this list likely indicates the concepts in which the students are most interested in the course. Analysis of student rating data indicates that each student rated on average 242 concepts (*mean* = 241.87, *std* = 132.12) for difficulty and 242 (*mean* = 242.97, *std* = 130.07) for relevance throughout

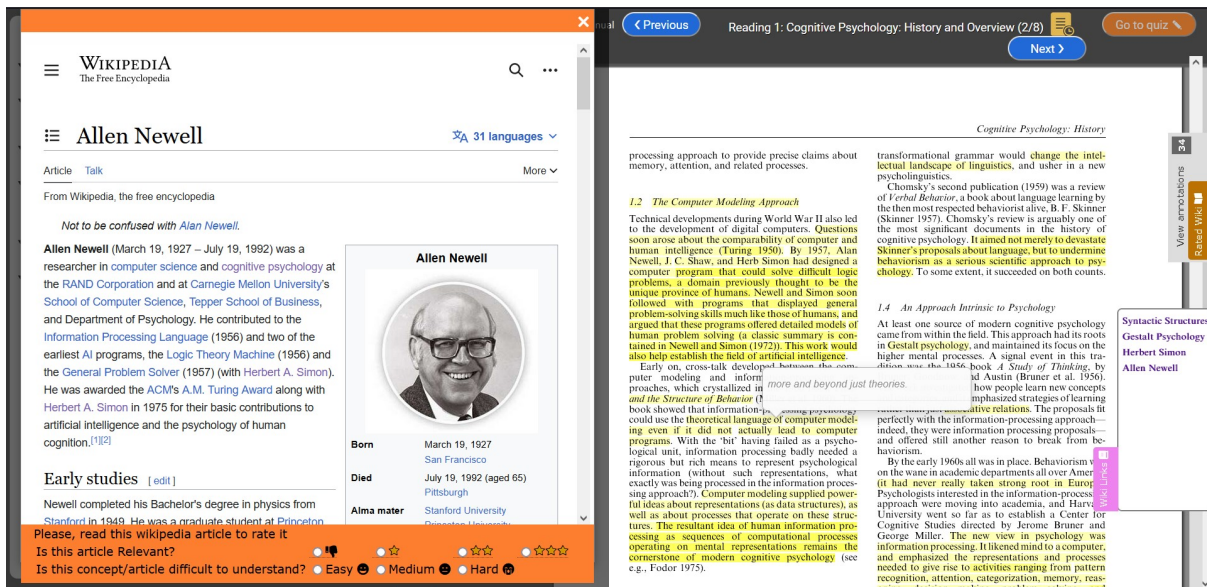


Figure 2: Once the student clicks on a link to a recommended Wikipedia article, it opens on the left side of the reading interface. The rating bar at the bottom allows the student to rate the relevance and difficulty of the recommended article.

Table 1: 10 Most Popular Wikipedia articles by number of students rating them as Relevant or Highly Relevant and as Medium or Hard Difficulty

Relevance	Difficulty
Change Blindness	Cognitive Science
Cognitive Science	Memory
Visual Perception	Change Blindness
Cognitive Psychology	Visual Perception
Saccade	Flicker
Experimental Psychology	Saccade
Cognitive Revolution	Cognitive Psychology
Iconic Memory	Distractions
Memory	Metadata
Hybrid Image	Mylifebits

the course duration. Note that it is considerably more than 110 ratings (10 per week) that the students were required to make to get the full score. This data indicates that the students were considerably engaged in examining and rating recommended Wikipedia articles.

The distribution of relevance and difficulty ratings for recommended articles rated is shown in Figure 3. As the data show, the majority of recommended articles were judged easy or medium difficulty by the class, although a noticeable number of articles were considered hard. From the prospect of relevance, the majority of articles were rated as relevant or highly relevant, although a good number were rated somewhat relevant and even not relevant.

To examine the articles rated as relevant or highly relevant, we counted the number of ratings for each of these articles (i.e., the number of students who rated this article as relevant or highly relevant) and plotted this data by ordering ar-

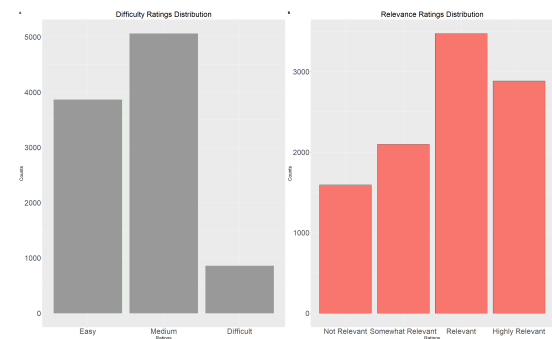


Figure 3: Distribution of Difficulty (left) and Relevance (right) ratings for recommended Wikipedia articles.

ticles by the number of ratings (Fig. 4). The data show that while a good number of concepts such as “Cognitive Science” and “Memory” were universally popular, approximately half of the relevant concepts such as “Probabilistic Reasoning” and “Knowledge Visualization” covered in Wikipedia articles were selected for examination by five or fewer students. This confirms our hypothesis that students in the same class have considerably different interests and opens up an opportunity for personalized rather than class-level recommendations.

As Fig. 3 shows, a considerable number of recommended Wikipedia articles were judged as not relevant. To understand how we can improve the recommendation process, we examined the concepts covered by these Wikipedia articles. The analysis revealed several problems. The dominant source of irrelevant recommendations was the PDF source of research articles. First, hyphenation frequently produces partial words such as “mecha” or “illus”, which sometimes have perfectly valid Wikipedia articles unrelated to the content of the course. Second, beyond their true con-




Figure 4: Relevant or highly relevant Wikipedia articles ranked by the number of ratings

tent, all articles have publication data, including named entities for publishers (“Princeton University Press”, “SAGE”, “IEEE”) and places of publication (“Hershey”, “Princeton”), which are usually present in Wikipedia. Another problem was the result of our attempt to recognize the names of researchers mentioned in the articles to offer students more information about them. Unfortunately, in a number of cases, these researchers were not prominent enough to appear in Wikipedia, while a different famous person with the same name was listed (i.e., “George Eyser”, “Terry Crews”), which resulted in referring to the wrong people. Finally, some perfectly valid concepts such as “priming” (in psychology) had different meanings in different areas and correspond to Wikipedia “disambiguation pages” with links to different meanings. Some students considered these pages irrelevant. The analysis demonstrated that most of the observed problems could be resolved by adding additional heuristics to our filtering process.

6. CONCLUSION

In this demo, we present a system that uses text mining to expand student reading options in graduate classes by recommending relevant Wikipedia articles for research papers assigned for mandatory reading. This approach enriches student course knowledge and allows students to personalize their readings by focusing on the most interesting concepts covered in the recommended articles. The system was used as a primary reading tool in a semester-long graduate course, enabling us to gain several interesting insights into student work with recommendations. In particular, we observed that about half of the articles rated as relevant or highly relevant were examined and rated by 5 or fewer students. It confirms that different students might be interested in different aspects of the course and opens opportunities for personalized recommendations. The current demo used a relatively simple text mining approach to extract interesting concepts mentioned in the text of the mandatory readings, yet the majority of recommended Wikipedia articles (and their concepts) were judged as relevant or highly relevant. The analysis of concepts judged as not relevant revealed several heuristics that could be used to improve our text-mining approach.

7. REFERENCES

- [1] R. Agrawal, S. Gollapudi, K. Kenthapadi, N. Srivastava, and R. Velu. Enriching textbooks through data mining. In *Proceedings of the First ACM Symposium on Computing for Development*, pages 1–9, 2010.
- [2] J. Barria-Pineda, P. Brusilovsky, and D. He. Reading mirror: Social navigation and social comparison for electronic textbooks. In *First Workshop on Intelligent Textbooks at 20th International Conference on Artificial Intelligence in Education (AIED 2019)*, volume 2225, pages 30–37. CEUR, 2019.
- [3] K. Bennani-Smires, C. Musat, A. Hossmann, M. Baeriswyl, and M. Jaggi. Simple unsupervised keyphrase extraction using sentence embeddings. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 221–229, Brussels, Belgium, Oct. 2018. Association for Computational Linguistics.
- [4] X. Liu and H. Jia. Answering academic questions for education by recommending cyberlearning resources. *Journal of the American Society for Information Science and Technology*, 64(8):1707–1722, 2013.