January 2023

Automatically Generating E-book Bookmarks from User Engagement

D Shin

Automatically Generating E-book Bookmarks from User Engagement

ABSTRACT

E-book reading applications and devices include functionality for bookmarking a page within an e-book to mimic physical bookmarks used with printed books. However, current e-book bookmarking features require the user to provide explicit manual input to bookmark a particular page and to learn and remember to use the bookmarking feature. This disclosure describes techniques for automatically bookmarking pages of potential interest within an e-book. With appropriate permissions from the user, page locations for the automatic bookmarks can be identified based on user engagement with the book content determined by inferring the reader's intent and attention when reading the e-book based on user-permitted interaction data. Parameters related to the user's engagement with the e-book content over time are encoded as a feature vector and converted to an adaptive histogram of attention scores. Additionally, the e-book contents are parsed with a suitable language model to obtain the e-book structure. The attention scores and e-book structure are combined by applying heuristics to generate high-quality bookmarks.

KEYWORDS

- E-book
- Digital book
- Digital bookmark
- Reading pattern
- Page turns
- Language model

BACKGROUND

Reading electronic books (e-books - books in a digital format) is popular. Users can read e-books on a variety of devices, such as smartphones, tablets, laptops, dedicated e-book readers, etc. E-book reading applications and devices include functionality for bookmarking a page within an e-book to mimic physical bookmarks used with printed books. Digital bookmarks enable users to mark pages of interest that they may wish to revisit at a later time. However, current e-book bookmarking features require the user to provide explicit manual input to bookmark a particular page and to learn and remember to use the bookmarking feature. This is burdensome and if a user doesn't use the feature, causes the user to waste time in searching for a page that the user forgot to bookmark.

DESCRIPTION

This disclosure describes techniques for automatically bookmarking pages of potential interest within an e-book. With appropriate permissions from the user, page locations for the automatic bookmarks can be identified based on user engagement with the book content determined by inferring the reader's intent and attention when reading the e-book based on user-permitted interaction data.

Selection of e-book pages for automatic bookmarking is performed using an algorithm that includes two main steps. The first step involves creating an adaptive histogram of the user's reading pattern for a given e-book. The second step involves parsing the attention scores represented by the adaptive histogram obtained in the previous step via a language model such as LaMDA [1] or another suitable model. The output of the model is discrete page stamps (converted from continuous attention scores) that correspond to the page locations that are candidates for automatic bookmarks.

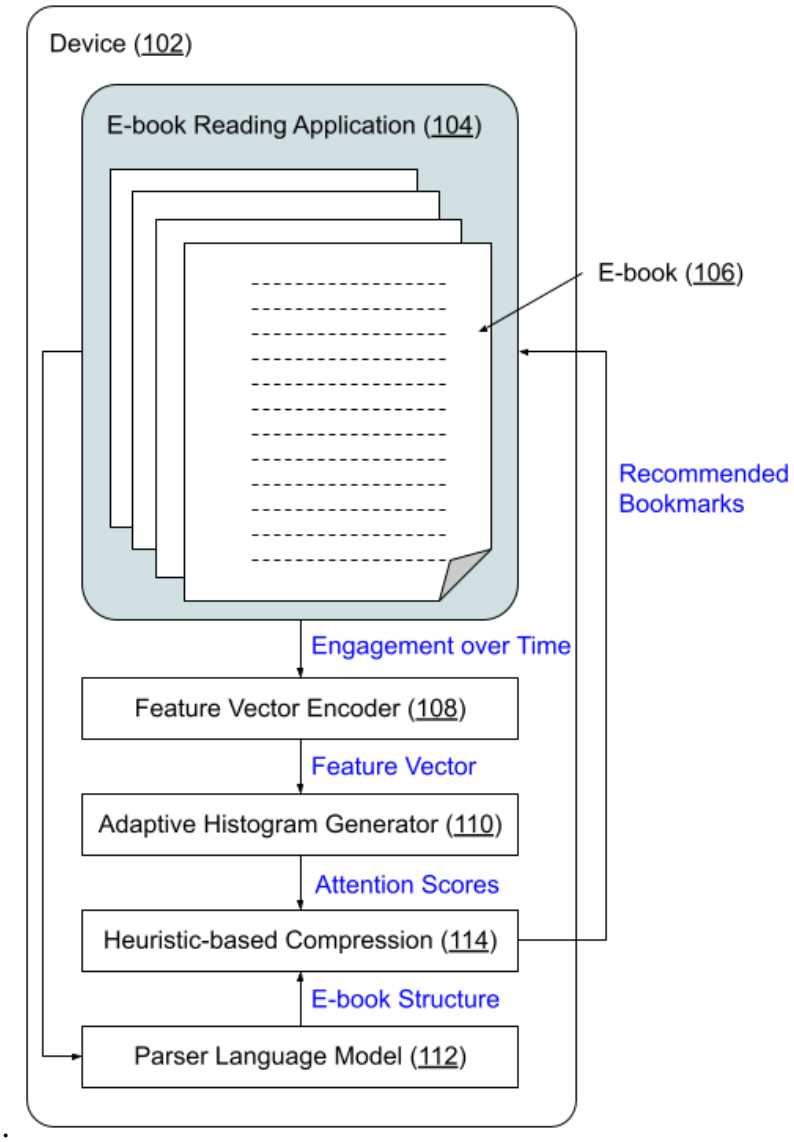

Fig. 1: Identifying e-book pages for bookmarking based on reading pattern

Fig. 1 shows an example of operational implementation of the techniques described in this disclosure. A user reads an e-book (106) within an e-book reading application (104) on a device (102). With the user's permission, parameters related to the user's engagement with the e-book content over time are encoded into a feature vector (108). An adaptive histogram is generated (110) to convert the information in the feature vector into raw scores that represent the user's attention with corresponding content within the e-book.

The e-book contents are parsed with a suitable language model parser (112) that outputs relevant information about the structure of the book, such as chapter boundaries, topics etc. The attention scores and e-book structure are analyzed by applying relevant heuristics (114) to generate a set of high-quality bookmarks that are recommended to the user for automated bookmarking, or based on corresponding user settings, automatically added to the e-book.

The feature vector can be generated by encoding the raw e-book content with user engagement as a one-dimensional temporal counter of the page state at discrete time intervals. For instance, each encoded value in the vector can signify the page number of the e-book on which the user spent the majority of that time interval. For example, with time discretization done at 30 seconds, a feature vector such as [1, 1, 1, 1, 1, 2, 2, 3, 3, 3, 3, 4, 4, 4, 4, 3, 3, 4, …] can indicate that the user spent 150 seconds (5 x 30s) on page 1, 60 seconds (2 x 30s) on page 2 etc. Moreover, the encoded information in the feature vector indicates dynamic navigation patterns while reading the e-book. For instance, the example feature vector above captures that the user moved back to page 3 after page 4 before proceeding forward again.

While a conventional histogram would count the number of state occurrences for each page within the feature vector, it would not capture the encoded navigation patterns between the pages. However, navigation patterns are an important indicator of user attention and engagement. For example, a user moving to a specific page rather than proceeding to the next page in sequence can signify higher user interest in the contents of that specific page.

To account for this, the feature vector of discrete page state times can be converted to an adaptive histogram that computes attention scores for each page by assigning higher weights to pages to which the user navigates out-of-sequence. For example, a conventional histogram for the feature vector [1, 1, 1, 1, 1, 2, 2, 3, 3, 3, 3, 4, 4, 4, 4, 3, 3, 4, …] will contain simple counts

for each page in the format page:count - [1:4; 2:2; 3:6; 4:5, …]. However, an adaptive histogram for the feature vector can be formed by assigning twice the weight to a page when a user moves to it out of sequence as [1:4; 2:2; 3:8 (4+2x2); 4:5; …]. In this case, the first four instances of page 3 are counted with the weight of 1 because they follow in sequence after page 2. However, the next two instances of page 3 are assigned double the weight because these occur out-of-sequence since the user moved back from page 4 rather than proceeding to page 5.

Pages with the highest scores in the adaptive histogram may not always align with the structure and the content of the e-book. For instance, pages with high attention scores can be located a couple of pages after the beginning of a chapter or the content of the pages may not be topically consistent. In the former case, it might be better to bookmark the page where the chapter begins whereas the latter situation may require finer grained bookmarking at the paragraph level within the page. Therefore, attention scores for each page in the adaptive histogram cannot be used directly to choose the candidate pages for bookmarking. Instead, candidate pages for bookmarking can be determined by analyzing the attention scores with input from a suitable language model that incorporates various aspects of the e-book content (e.g., structure, topics, semantics, etc.) at appropriate resolution (e.g., sections, chapters, pages, paragraphs, etc.).

The combined analysis of the attention scores from the adaptive histogram and the output of the language model regarding e-book content can be performed by employing a heuristic to find multimodal peaks in the histogram with any suitable approach, such as a single dimensional Gaussian-mixture fit. The means or modes of each Gaussian component can correspond approximately to a page suitable for bookmarking. Each mean or mode of the Gaussian fit can then be adjusted based on information about the e-book contents as indicated by the output of the

language model. For example, if the initial recommended candidate bookmark based on the adaptive histogram scores is at page 3 and the output of the language model indicates that pages 2-10 are about the same topic, then the bookmark recommendation can be adjusted to page 2 instead of page 3. Such adjustments to the initial approximate candidate bookmarks to account for the e-book contents can be achieved with L2 minimization in the page numbering metrics for each histogram value against the entire result set obtained via the language model.

The described techniques can be implemented for automatic bookmark generation for any device, platform, or application that supports reading e-book content. Settings and preferences related to automatic bookmarks can be specified by the users. While all components required for generating the automatic bookmarks are located locally on the user device in the operational implementation described above, with user permission, one or more of the components can be located externally, such as on a server, cloud-based platform, etc. Implementation of the techniques can substantially reduce the inconvenience and burden of manually bookmarking e-book pages, thus enhancing the user experience (UX) of reading e-books and revisiting interesting and important content within them. While the techniques are described with reference to e-books, automatic bookmarks can also be provided for other types of electronic documents such as reports, presentations, long-form articles, etc.

Further to the descriptions above, a user may be provided with controls allowing the user to make an election as to both if and when systems, programs or features described herein may enable collection of user information (e.g., information about a user's electronic books or documents, social network, social actions or activities, profession, a user's preferences, or a user's current location), and if the user is sent content or communications from a server. In addition, certain data may be treated in one or more ways before it is stored or used, so that

personally identifiable information is removed. For example, a user's identity may be treated so that no personally identifiable information can be determined for the user, or a user's geographic location may be generalized where location information is obtained (such as to a city, ZIP code, or state level), so that a particular location of a user cannot be determined. Thus, the user may have control over what information is collected about the user, how that information is used, and what information is provided to the user.

CONCLUSION

This disclosure describes techniques for automatically bookmarking pages of potential interest within an e-book. With appropriate permissions from the user, page locations for the automatic bookmarks can be identified based on user engagement with the book content determined by inferring the reader's intent and attention when reading the e-book based on user-permitted interaction data. Parameters related to the user's engagement with the e-book content over time are encoded as a feature vector and converted to an adaptive histogram of attention scores. Additionally, the e-book contents are parsed with a suitable language model to obtain the e-book structure. The attention scores and e-book structure are combined by applying heuristics to generate high-quality bookmarks.

REFERENCES

1. "LaMDA: our breakthrough conversation technology" available online at https://blog.google/technology/ai/lamda/ accessed January 20, 2023.
2. Bodapati, Sravan Babu, Sriraghavendra Ramaswamy, and Gururaj Narayanan. "A Machine Learning Approach to Detecting Start Reading Location of eBooks." In *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 1522-1529. IEEE, 2018.