



## **A Visual Approach for Health Information Exploration: Adaptive Levels of Visual Granularity and Interaction Analysis**


### **Stefan Lengauer**

(Institute of Visual Computing, Graz University of Technology, Graz, Austria  
 <https://orcid.org/0000-0001-5136-4320>, [s.lengauer@tugraz.at](mailto:s.lengauer@tugraz.at))


### **Lin Shao**

(Fraunhofer Austria Center for Data Driven Design, Graz, Austria  
 <https://orcid.org/0000-0002-5118-9515>, [lin.shao@tugraz.at](mailto:lin.shao@tugraz.at))


### **Hossein Miri**

(International School of Engineering, Chulalongkorn University, Bangkok, Thailand  
 <https://orcid.org/0009-0007-2522-1888>, [hossein.m@chula.ac.th](mailto:hossein.m@chula.ac.th))

### **Michael Bedek**

(Institute of Psychology, University of Graz, Graz, Austria  
 <https://orcid.org/0009-0005-1397-5186>, [michael.bedek@uni-graz.at](mailto:michael.bedek@uni-graz.at))


### **Cordula Kupfer**

(Institute of Psychology, University of Graz, Graz, Austria  
 <https://orcid.org/0009-0008-5955-4725>, [cordula.kupfer@uni-graz.at](mailto:cordula.kupfer@uni-graz.at))


### **Maria Zangl**

(Institute of Psychology, University of Graz, Graz, Austria  
 <https://orcid.org/0009-0000-8388-5518>, [maria.zangl@uni-graz.at](mailto:maria.zangl@uni-graz.at))


### **Bettina Kubicek**

(Institute of Psychology, University of Graz, Graz, Austria  
 <https://orcid.org/0000-0003-0725-1657>, [bettina.kubicek@uni-graz.at](mailto:bettina.kubicek@uni-graz.at))

### **Barbara Dienstbier**


(Institute of General Practice and Evidence-based Health Services Research,  
Medical University of Graz, Graz, Austria  
 <https://orcid.org/0009-0001-4631-3729>, [babsi.dienstbier@gmx.at](mailto:babsi.dienstbier@gmx.at))

### **Klaus Jeitler**

(Institute of General Practice and Evidence-based Health Services Research,  
Medical University of Graz, Graz, Austria  
Institute for Medical Informatics, Statistics and Documentation,  
Medical University of Graz, Graz, Austria  
 <https://orcid.org/0000-0002-1520-8061>, [klaus.jeitler@medunigraz.at](mailto:klaus.jeitler@medunigraz.at))


### **Cornelia Krenn**

(Institute of General Practice and Evidence-based Health Services Research,  
Medical University of Graz, Graz, Austria)

 <https://orcid.org/0000-0001-9292-4291>, cornelia.krenn@medunigraz.at)


### **Thomas Semlitsch**

(Institute of General Practice and Evidence-based Health Services Research,  
Medical University of Graz, Graz, Austria)

 <https://orcid.org/0000-0002-1763-7855>, thomas.semlitsch@medunigraz.at)


### **Carolin Zipp**

(Institute of General Practice and Evidence-based Health Services Research,  
Medical University of Graz, Graz, Austria)

 <https://orcid.org/0000-0002-1587-1269>, carolin.zipp@medunigraz.at)


### **Dietrich Albert**

(Institute of Psychology, University of Graz, Graz, Austria)

 <https://orcid.org/0000-0001-6701-3058>, dietrich.albert@uni-graz.at)


### **Andrea Siebenhofer**

(Institute of General Practice and Evidence-based Health Services Research,  
Medical University of Graz, Graz, Austria)

Institute for General Practice, Goethe University Frankfurt am Main  
 <https://orcid.org/0000-0002-6980-2103>, andrea.siebenhofer@medunigraz.at)

### **Tobias Schreck**

(Institute of Visual Computing, Graz University of Technology, Graz, Austria)

 <https://orcid.org/0000-0003-0778-8665>, tobias.schreck@tugraz.at)

**Abstract:** The effective and targeted provision of health information to consumers, specifically tailored to their needs and preferences, is indispensable in healthcare. With access to appropriate health information and adequate understanding, consumers are more likely to make informed and healthy decisions, become more proficient in recognizing symptoms, and potentially experience improvements in the prevention or treatment of their medical conditions. Most of today's health information, however, is provided in the form of static documents. In this paper, we present a novel and innovative *visual health information system* based on adaptive document visualizations. Depending on the users' information needs and preferences, the system can display its content with document visualization techniques at different levels of detail, aggregation, and visual granularity. Users can navigate using content organization along sections or automatically computed topics, and choose abstractions from full texts to word clouds. Our first contribution is a formative user study which demonstrated that the implemented document visualizations offer several advantages over traditional forms of document exploration. Informed from that, we identified a number of crucial aspects for further system development. Our second contribution is the introduction of an *interaction provenance visualization* which allows users to inspect which content, in which representation, and in which order has been received. We show how this allows to analyze different document exploration and navigation patterns, useful for automatic adaptation and recommendation functions. We also define a baseline taxonomy for adapting the document presentations which can, in principle, be leveraged by the observed user patterns. The interaction provenance view, furthermore, allows users to reflect on their exploration and inform future usage of the system.

**Keywords:** Document Visualization, Adaptive Visualization, Interaction Analysis, Health Information, Evaluation

**Categories:** H.3.5, H.3.7, H.4.0, H.5.1

**DOI:** 10.3897/jucs.150679

## 1 Introduction

Consumer Health Information Systems (CHISs) have become an important tool in modern healthcare and serve a variety of functions. They can provide consumers with a comprehensive understanding of a disease, specifically addressing general knowledge of health-related issues, effects, ways, and measures to maintain and possibly restore health. They can also enable early detection, diagnosis, treatment, palliation, rehabilitation, and follow-up care for diseases, along with associated medical decisions, care, and coping strategies for daily life with these diseases [Arbeitsgruppe GPGI 2016]. Typically, CHISs are provided in a static and linear manner, meaning the same medical content is presented to everyone in the same structure. However, a linear reading or navigation may not be the best solution for everyone to extract relevant information because patients differ in terms of their prior knowledge, information needs, personal preferences and styles, and health situation, which may depend on factors such as gender, age, personality, and perception [Bunge et al. 2010]. To address this problem, an *adaptive and interactive visual CHIS that supports document exploration with adaptive focus and detail views* is needed. By tailoring content to individual user needs and preferences, such as adjusting the level of detail or focusing on specific topics, this type of system would allow users to personalize their experience.

The primary research goal of this work is to develop novel concepts for advanced, interactive, adaptive, and visual CHIS (called Adaptive CHIS (A<sup>+</sup>CHIS)). We selected Type 2 Diabetes Mellitus (T2DM) as a pilot disease for our research because it is a complex disease with high prevalence and relevance to the public health system. Also, its progression over time requires changes and adaptations, including tailored knowledge, flexible treatment, new drug classes, improved patient education, sustainable follow-up practices, and screening for complications. Managing T2DM is still a difficult and time-consuming task, as it is common, serious, and under-treated. This is a major challenge for healthcare services, and patients and therapists must be prepared to deal with it effectively. Consequently, those affected by T2DM have an ongoing requirement for timely and pertinent information [Standl et al. 2019].

In this paper, we present a novel visual document exploration system with multi-dimensional adaptivity to help health information consumers better understand medical content by combining close and distant reading approaches. It is targeted towards non-medical users of all adult age groups, which have either T2DM, are a relative to a T2DM patient or are interested in the disease for another reason. After a development and evaluation phase the A<sup>+</sup>CHIS should eventually be freely discoverable on the Internet. We work thus with an unpremeditated userbase in mind, which should be able to use the CHIS intuitively, without the need for a supervised introduction.

To this end, we propose an innovative document exploration system that provides multi-level navigation from high-level (topic overview) to mid-level (keyword occurrence and highlighting) to low-level (full text) views (Figure 1). The basic idea is to allow users to efficiently navigate through documents, overview the content, find topics of interest, and finally switch to a close read on specific information contents. In our system, we make use of well-known document visualization approaches to enhance the learning process of medical content. To visualize high-level structures of a document, we propose dynamic table-of-content which represents sub-chapters by means of a Word Cloud containing keywords from a topic-modeling approach. High-level structures and mid-level document information are linked by using Tile Bars as visual navigator. The visual navigator shows topic occurrences within the underlying document and allows users to quickly explore the content by text snippets. Although these text visualization techniques

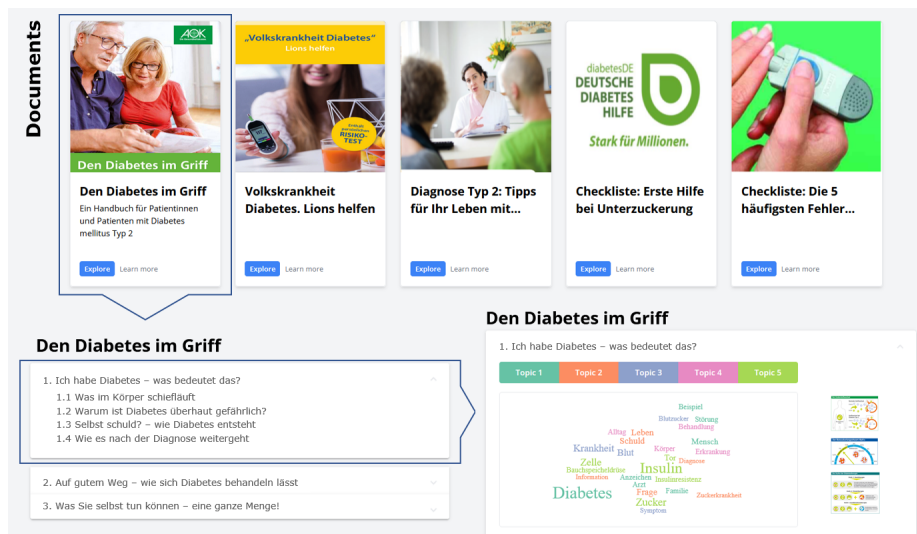


Figure 1: The proposed A<sup>+</sup>CHIS supports dynamic levels of detail. A Document Library (top) allows users to select one particular document they want to explore. At document level, an interactive Table of Contents (bottom left) preserves the global linear structure of the document while a chapter's substructure and content is visualized with dedicated visualization techniques for textual and pictorial data.

are not novel themselves, we tie them together in an integrated implementation prototype and adapt them for the specific requirements of the health domain. The system also serves as a platform to test and evaluate approaches for adaptive document and interaction provenance visualization (cf. also Section 5). Our system introduces the notions of level of detail and adaptive visual presentations for document-based health information exploration for T2DM. Existing CHIS are largely static in nature and do not adapt either of these dimensions to their users.

As a consequence, we conducted a user study to characterize the usage behavior of health information seekers adopting our approach. We show the usability of our system by comparing linear reading with our multi-level approach, and illustrate the results by two provenance visualizations.

This paper presents a comprehensive extension of one of our previous papers [Shao et al. 2023] delving further into our developed system and offering additional insights and analyses. Besides a more verbose and detailed outline of our work, this paper comes with the following additions with regards to its predecessor:

1. An extension to our system design with an updated visual representation as well as newly-added components such as a Document Library for the exploration at documents level and an alternative to the Word Cloud representation (Topiccloud (ToC));
2. A thorough analysis of the interaction data obtained through supervised evaluations with actual users;

3. Different customized visual analytics tools for analysing and evaluating said interaction data.

In the next section (Section 2) we provide an overview on relevant previous works from the consumer health information domain and reveal the research gap regarding adaptive systems for said domain, before our proposed system design is outlined in Section 3. Section 4 describes a formative evaluation that incorporates quantitative and qualitative methods, performed by a number of participants on an implementation prototype. Following this formative evaluation it was observed that the rich interaction data collected in its course merits the effort of a more detailed analysis and the development of additional visual analytics tool (Section 5). On that note, we also investigated the suitability of user interactions for proposing adaptive visualizations. Section 6 and 7 conclude the paper with a thorough discussion regarding next steps and open research challenges.

## 2 Related Work

In this section, we present an overview of important visualization techniques that have inspired our approach and review the need for adaptive and interactive consumer health information systems.

### 2.1 Visualization Techniques for Text Documents and Health Care

Visualizing large text corpora is a challenging task. Usually, the involved data sets are inherently complex, containing structural and content-related information. Most linguistic and text visualization approaches rely on text-mining techniques to reveal semantic information from raw text data. Therefore, simple statistical processing (e.g., word frequency and bag-of-words concept) as well as natural language processing approaches (e.g., named-entity recognition, relationship extraction and sentiment analysis) may be used [Alencar et al. 2012; Kucher & Kerren 2015].

A widely-used visualization technique for text data is the Word Cloud representation (also known as Tag Cloud) which presents an overview of the most frequent or important words by using different type or font sizes [Heimerl et al. 2014]. This technique is also known as distant-reading technique [Moretti 2005] and allows users to approach literature in a new way. Instead of reading texts in the traditional way, i.e., linear or close-reading, the focus of distant-reading approaches is to count, graph, and map textual data by a visual representation [Jänicke et al. 2015]. In recent years, much research has been conducted on distant-reading and Word Cloud visualizations. For instance, Kim et al. [Kim et al. 2011] proposed WordBridge, which utilizes graph-based visualization techniques to connect multiple Word Clouds with information-rich edges. Further extensions of Word Cloud exist that focus on semantic contour lines [Wu et al. 2011] and images [Gu et al. 2017]. In our work, we rely on traditional Word Clouds to foster distant-reading within single documents.

For the exploration of larger document collections, additional document features such as metadata information and co-authorships, could be considered to gain a better understanding of the contents of those documents [Görg et al. 2013; Isenberg et al. 2017]. Another interesting approach by Strobelt et al. [Strobelt et al. 2009] called Document Cards, utilizes a mixture of images and important keywords to visualize key semantics of a document. To visualize distributional properties within a document, Tile Bars [Hearst 1995; Keim & Oelke 2007] could be considered, which is a compact pixel-based

visualization technique that reveals the relative length of a document and the relative frequency of one or more query terms. In our work, we utilize Tile Bars to represent the relative frequency and distribution of terms from a Word Cloud.

Data visualizations are becoming increasingly important for various fields of application, as well as in healthcare. Visual representations may help patients as well as physicians to gain a better understanding of health records, e.g., information on medical diagnostics, treatments, and health histories [Rind et al. 2013]. For example, the LifeLine system was among the first exploration systems that supports visual patient treatment histories [Plaisant et al. 1998]. An extensive survey about visualization techniques for electronic health records and population health records are given in [Wang & Laramée 2022].

Recently, many of the mentioned document visualization techniques are also applied in a medical context. For instance, Facetatlas by Cao et al. [Cao et al. 2010] used linked Word Clouds to visualize entity-relational text document of diseases such as Type 1 and Type 2 Diabetes Mellitus. The linked Word Clouds are used to represent global relations by using a density map and local relations by using edge bundling techniques. Another interesting multifaceted text visualization is SolarMap [2011] which combines a labeled contour-based cluster visualization with a radially-oriented word cloud. Furthermore, SolarMap can visualize topic distribution of entities from one facet together with keyword distributions that convey the semantic definition along a secondary facet.

With the advent of novel visualization techniques in different domains, visualization literacy, i.e., user understanding and discovery of visual patterns, is becoming increasingly important. Developing visual literacy is essential to support cognition and evolve toward a more informed society [Firat et al. 2022]. In our work, we intend to increase visual and health literacy by gathering user information during exploration and providing adapted health information based on that.

## 2.2 Need for Adaptive CHISs

As part of this work, we examined current sources of CHISs related to T2DM across multiple media platforms, including websites, digital documents, print media, apps, and videos. Our goal was to identify elements and modes of presentation within a representative sample of these sources that users can customize to their needs and preferences. Our results suggest that the potential for adaptation in existing CHIS is only realized to a limited extent. We did not find any adaptive elements in print media or digital documents while websites, apps, and videos offer some customization options related to presentation format, such as adjusting font size and color. Some CHISs also included features such as text-to-speech or language-switching [Bundesministerium für Gesundheit 2022; Government of Canada 2022]. However, in terms of personalized medical information, only a few CHISs had mechanisms to pre-filter medical content based on a user's diabetes profile [British Diabetic Association 2022]. Most CHISs included a standard table of contents, with or without hyperlinks to the respective chapters. Some sources also contained links within the text or cross-references to other sections or chapters. However, none of the T2DM CHISs we analyzed used a visual document exploration system with multi-dimensional adaptivity for health information consumers.

These results show that existing CHISs on T2DM fall short of the potential of presenting health information in an *interactive*, *adaptive* and/or *personalized* way, while there is evidently a need for it. The knowledge domain of T2DM is complex and comprehensive, with a wide range of information sources (brochures, websites, medical doctors, etc.) and high diversity of topics (such as symptoms, treatments, nutrition, etc.). This

might be overwhelming for laypersons without medical expertise seeking knowledge in the field. Such complex information situations usually put a high *intrinsic cognitive load* [Sweller 2005] on the working memory during information processing and often lead to information seekers applying heuristics and cognitive biases at every stage of information processing. Such cognitive biases, misconceptions, and even myths about T2DM may lead to unhealthy behavior with severe health-related consequences.

An *interactive* CHIS has the potential to (i) track behavioral patterns and explicit feedback of consumers, (ii) interpret these indicators in terms of certain cognitive biases (e.g., the confirmation bias), and (iii) intervene if necessary (e.g., by suggesting other pieces or sources of information).

An *adaptive* CHIS can match the information units to the users and their current information needs. It can thus balance the *intrinsic cognitive load* to a medium level and ensure that the consumer is neither too bored nor too overwhelmed. This is in line with the transfer of Vygotsky's concept of the *zone of proximal development* to digital learning environments [Luckin 2010], the outer fringe as suggested by the competence-based knowledge space theory [Heller et al. 2006] and constitutes a solid basis for an enjoyable flow for consumers [Schiefele & Raabe 2011]. All these theories emphasize that a medium difficulty of information units lead to a successful processing outcome. The intermediate goal is to successively reach more advanced levels of learning outcomes as suggested by Krathwohl [Krathwohl 2002] which means not simply remembering information, but also applying and evaluating it.

A *personalized* CHIS can foster consumers' personal commitment to engage with the system and information, and thus help to close the 'intention-behaviour gap' or 'attitude-action gap' [Schwarzer 2008]. This is the ultimate goal of any CHIS. We strive to achieve this through different added values of our advanced CHIS compared to more 'traditional' digital CHISs (e.g., a brochure in PDF format or plain webpage): the guarantee of high quality and evidence-based medical information, the reduction of complexity to a medium level, and the recommendation of information units that fit a consumers' information needs. In addition, tools and functionalities to get an overview of the knowledge domain, to efficiently answer certain questions, and to easily navigate through different sub-topics will be offered for good user experience. With our formative evaluation activities (Section 4) we monitor progress towards these goals.

### 3 Adaptive Document Exploration Design

In the following Section 3.1 we discuss the design considerations, taking into account the intended userbase and use case scenarios. Informed by this, the resulting system design is outlined in detail in Section 3.2.

#### 3.1 Design Requirements/Considerations

In view of our intended users, who consists of all adult age groups, genders, and digital proficiency we decided to rely on document visualization techniques that are intuitive to understand and easy to use. Although the participants of the formative evaluation study which was conducted within the scope of this project (Section 4) received a brief introduction into visual components of the system, we want to support the scenario that an uninformed user, which discovers the CHIS on the Internet is able to use it without any prior supervised introduction. That is, visualizations should base upon tried and tested concepts, which have been shown to be usable by the vast majority of people. Yet,

at the same time, more advanced users should have to option to alternatively work with more complex visualizations, which allow for a more efficient exploration at the cost of a steeper learning curve. prior

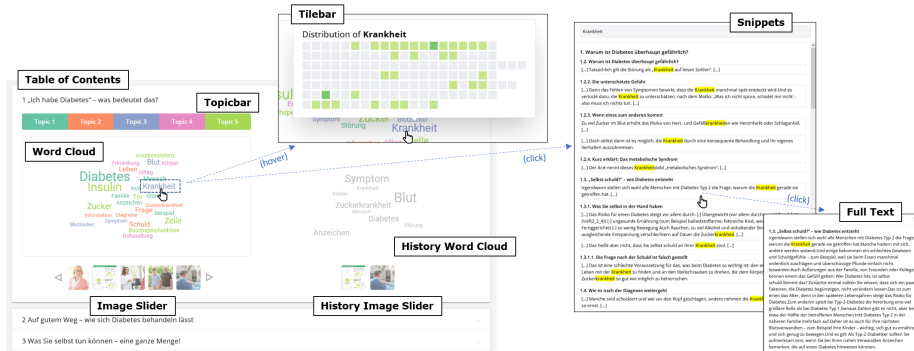


Figure 2: The main components of  $A^+$ CHIS shown by an example of exploring a German diabetes health brochure [2021]: Table of Contents (ToC), Word Cloud (WC)/History WC (HWC), Image Slider (IS), Tilebar (TiB), Topicbar (ToB), Snippets (Snps), and Fulltext (FullT). Different actions (illustrated as blue arrows) allow a user to navigate from one view to another.

### 3.2 Implementation

Informed by the considerations above we designed a system comprised of several components, which support the exploration of documents at different levels of visual granularity. Starting from a Document Library (DL) showing an overview of the documents supported by the  $A^+$ CHIS, a user is able to explore individual documents on both high and low levels (Figure 2). For the prior, the system provides an expandable and interactive Table of Contents (ToC) while the latter is enabled through a series of text abstraction methods such as Word Cloud (WC), fingerprinting in the form of a Tilebar (TiB), and topic modeling by means of a Topicbar (ToB). Pictorial content is provided in an Image Slider (IS) component. On the lowest level, a user can also review sections of the original text sources in the form of Snippets (Snps) or even the untampered full text. These different concepts are implemented in different subsystems such that an uninterrupted exploration process is possible while alternating between levels. In the background, we track user interactions to determine which parts of the content have already been visited and consumed by the user. This information is also displayed to the user in order to indicate which information has not yet been scrutinized.

For the technical realization, we rely on the *Angular*<sup>1</sup> framework for the implementation of the different components in the frontend, backed up by a *Django*<sup>2</sup>-driven backend, connected to a Postgres database for storing interaction data. This scalable setup allows for many users operating the system concurrently through the web.

<sup>1</sup> <https://angular.dev/>

<sup>2</sup> <https://www.djangoproject.com/>

### *Document Library*

The DL is the first view a user is greeted with upon logging into the A<sup>+</sup>CHIS. All available documents are presented in a grid arrangement with their respective covers and titles (Figure 1). Upon hovering over a certain document a preview appears, showing both related metadata and a histogram of the document's most frequent terms. This initial view should serve the users in determining which of the documents is the most appropriate for addressing their respective information need. Clicking on a document transfers the user to its ToC.

### *Table of Contents*

The exploration process on document level is supported by an interactive ToC. We base this view on a document's inherent linear structure with chapters, sub-chapters, and so on. While we aim to preserve the outermost structure ('chapters' in most cases) we abstract all lower levels of structure and content with dedicated text- and multimedia visualizations (Figure 1). To this end, we employ WCs – and most recently also an alternative representation (Section 6.1) – for the textual content, ISs for the pictorial content, and further subsystems (TiB and Snps) for structure and lower levels of visual granularity. These dedicated visualizations are interweaved into the linear chapter structure, based loosely on the *document card* design concept by Strobelt et al. [Strobelt et al. 2009] – i.e., different visualization techniques are used to display the textual and visual contents. On a per-chapter basis, these visualizations can be expanded or collapsed.

Additionally, already 'consumed' content is tracked and indicated with a 'history' version of the respective WCs and ISs. Specifically, terms and images are added to these components after they have been reviewed (i.e., clicked on) by the user. This History WC (HWC) visualizes the context of the exploration to the user. Alternatively, it can be used to display non-clicked terms as to suggest content to the user.

### *Word Cloud with Topicbar*

To generate the word clouds, natural language processing is used to extract 'significant' terms from each chapter. This pre-processing comprises steps such as tokenization (separation/segmentation into individual parts) and stop-word removal (filtering of irrelevant/insignificant words). Subsequently, the set of remaining words are subjected to a lemmatization (transformation into their canonical form or dictionary form) and grammatically tagged (part-of-speech tagging). However, we do not use the resulting normalized words directly to fill the WC, but instead subject them to the Latent Dirichlet Allocation [Blei et al. 2003] in order to obtain meaningful topic models on them. The appropriate number of topic models per-chapter (5) was determined empirically by visually evaluating the resulting topics for different values in the low integer range. That is, each topic model is defined by a weighted vector containing a subset of the chapter's extracted nouns.

The concated vectors of all topics serve as the input for a chapter's WC, where the weights are used to determine a word's size within the cloud. For the arrangement of terms, we rely on the *Wordle word cloud* algorithm [Steele & Iliinsky 2010] while the mapping between words and their belonging topic is established through a qualitative color mapping (Figure 2, WC). Note that this means certain terms could appear redundantly as topics can exhibit overlapping term compositions.

Since the concurrent display of all topics can be overwhelming, we provide a means to toggle the visibility of individual topics. This subsystem – the ToB directly above

the WC – consists of 5 colorized toggle buttons mapping to the respective topic models. Hovering over a term toggles the Tilebar component for the respective term above it, while clicking it initiates the Snps view. To this end, we track the interactions – how often the user has clicked a certain term – in the WC. These click counts are the basis for the so-called History WC (HWC) on the right-hand side of a chapter visualization (Figure 2, HWC) where the count determines the size of a term in the cloud. The same hover-and-click interactions as with the ‘regular’ word cloud are possible with the HWC.

#### *Tilebar*

Hovering over a term in the WC triggers the display of the TiB component above it, which allows the user to efficiently grasp the term’s occurrences over the whole document. This visualization is inspired by the *literature fingerprinting* concept by Keim and Oelke [Keim & Oelke 2007] which shows various document properties in a drilled-down manner. To this end, we compute the respective term’s frequency over equal-sized text chunks and visualize the resulting ‘intensities’ in a colorized fashion, following the linear document structure (i.e., from top to bottom and left to right, see Figure 2, TiB). That is, the rows of the grid symbolize the document’s chapters and the columns the ordered text chunks within a chapter. Cells which stand for text chunks in which the term does not appear at all are filled with uniform gray color. The TiB allows a user to quickly answer such questions as “*does another chapter also cover this topic?*” or “*how frequently is it mentioned overall?*”.

#### *Snippets/Fulltext*

The above mentioned abstractions are vital to gain an overview of the information covered by a document and determine which sections are the most appropriate to answer a specific information need. Ultimately however, it is necessary to provide a user with text chunks to enable them to answer their specific information need. To address this issue, we added two additional levels of visual granularity. Firstly, a *Snippets* view which pops up to the right hand-side of ToC if a term in the WC or HWC is clicked. Within this view, all sentences containing the clicked term are displayed with a highlighting of the term (Figure 2, Snps). Handles at the beginning and the end of a sentence allow to reveal the preceding and succeeding sentence. Those can be clicked iteratively to display larger parts of the document before and after the found position. Alternatively, the section headers, which are also shown in the snippets view, can be clicked to display a section’s whole content immediately (Figure 2, FullT). The top of the Snps view also contains a Searchbar, which allows to readily change the term in question.

#### *Image Slider*

Besides the abstractions for textual content (WC/HWC), we would like to indicate the presence of a document’s pictorial content. To this end, we employed an off-the-shelf IS component, next to the WC to display a chapter’s images (Figure 2, IS). The thumbnails of the images can be viewed directly in this slider, yet if further details are needed, an image can be clicked which shows it together with its caption in a full screen modal dialog (henceforth referred to as IS<sub>s</sub> and IS<sub>l</sub> respectively). As this component shows just five images at a time, we aim to determine an image’s relevance in order to sort the list of images, resulting in the most relevant images being shown initially. To determine said relevance, we make two assumptions. Firstly, we assume that images without a caption (e.g., scenic backgrounds at the beginning of a chapter) are rather unimportant. Secondly,

we split the set of images with captions into two tiers with the first tier being comprised of images showing tables, diagrams, and flow charts, or convey any sort of structured information, while all the others belong to the second tier. The information whether or not an image has a caption is obtained in the extraction process. In a similar fashion to the provenance version of the WC – the HWC – we provide an additional image slider in the bottom of the HWC, showing exclusively the chapter’s images which have already been clicked (and thus ‘consumed’) by the user.

## 4 Formative Evaluation

The formative study aimed to (a) explore how the design of the system and its components would be perceived by users, (b) evaluate the system in comparison to a linear and static CHIS, (c) highlight potential areas for improvement, and (d) identify prospective research questions. To maintain methodological consistency and participant comparability, we focused on one particular document for the evaluation setup. That is, we used a T2DM information brochure of the German health insurance provider AOK [Baumgart et al. 2021] as data basis for evaluation. The text document in PDF format comprises more than 130 pages of comprehensive and detailed health information, including figures, tables, and info-graphics. Full texts with health information were extracted with Adobe PDFBox<sup>3</sup> library and the images were extracted manually. Subsequently, the sub-chapters and images were sorted by the main chapters.

### 4.1 Participants

Overall, 12 participants (four females) took part in the study, which seems quite small for an empirical investigation, but is reasonably large enough for a formative evaluation, which is intended to gather initial feedback, in particular qualitative one. A balanced gender ratio was not primarily pursued, but we aimed to include various potential users of A<sup>+</sup>CHIS resulting in a sample with ages ranging from 26 and 62 years ( $M = 40$  yrs.,  $SD = 14$  yrs.) as well as different levels of knowledge and competence. On a 5-point rating scale (from 0 := very low to 4 := very high) they self-assessed their prior knowledge of T2DM ( $M = 1.00$ ,  $SD = 1.21$ ), computer and software skills ( $M = 2.25$ ,  $SD = .97$ ), as well as previous experiences with visualizations ( $M = 2.58$ ,  $SD = 1.24$ ).

### 4.2 Procedure

The overall procedure for the participants can be divided into the following phases: (i) *Instruction*, (ii) *Cognitive Walkthrough*, (iii) *Forced Choice* and finally, (iv) *Semi-Structured Interviews*. The phases (ii) to (iv) are described in more detail in according subsections below. The sessions that lasted between 60 and 90 minutes per participant were carried out individually, lead by one investigator. Three of the authors took the role of an investigator, each for four participants. In phase (i), i.e. *Instruction*, participants received a short explanation of basic CHIS functions, such as search functions, to ensure that they all began the evaluation process from a common usage knowledge base. The participants were then set in a real-world usage scenario: they were asked to imagine that they themselves or someone in their family was diagnosed with T2DM during a health check-up. That is why they want to find out more about this disease.

<sup>3</sup> <https://pdfbox.apache.org/>

The above-mentioned AOK brochure [2021] was presented in A<sup>+</sup>CHIS, where participants started the first task with the ToC view of the brochure, as well as in a PDF viewer (Adobe Acrobat Reader). The comparison with the PDF viewer was chosen not only because this is the original, static format of the brochure, but also because PDF viewers are widely used and therefore people are usually accustomed to them. The PDF viewer is therefore a challenging benchmark that allowed us to evaluate the expected added value of A<sup>+</sup>CHIS compared to conventional CHISs. The investigators recorded all audio and on-screen activities and additionally made manual notes of their observations.

### *Cognitive Walkthrough*

To encourage interaction with A<sup>+</sup>CHIS, showing how intuitive its functions are and how quickly the content can be grasped, participants were given pre-defined tasks to explore A<sup>+</sup>CHIS. This evaluation method is known as a Cognitive Walkthrough (CWT) [Hollingsed & Novick 2007]. The pre-defined tasks also enabled comparable conditions across all participants for the subsequent evaluation steps. We defined them in such a way that they represent realistic search and evaluation tasks in the course of an information search. For the purpose of the study, the tasks also had to be linked to a measurable goal achievement, for example, the participants had to find and report a specific piece of information. It had to be possible to achieve this goal in both PDF viewer and A<sup>+</sup>CHIS. For each tool, approximately the same number of tasks were designed, including tasks from the two categories of ‘generating an overview’ vs. ‘finding specific information’. Initially, we defined 11 tasks: four were aimed at using the Word Cloud including Topicbar (WC1–4), four at using the Tilebar (TiB1–4), and three at using the Image Slider (IS1–3). Since we intended to compare the PDF viewer with A<sup>+</sup>CHIS, two parallel versions were created for the 11 tasks that were comparable in terms of difficulty and functions used. For example, the task TiB4 ‘In which chapter would you most likely start if you wanted to find out more about blood pressure?’ had the parallel version ‘In which chapter would you most likely start if you wanted to find out more about insulin?’. This resulted in a total of 22 tasks (Table 1) that each participant processed in a within-subject design via the PDF viewer and A<sup>+</sup>CHIS. We balanced task order in terms of system and components to avoid sequential effects. In addition to the CWT, participants were asked to express their thoughts during the tasks (i.e., *think-aloud*). The investigators also noted behavioral observations during the CWT tasks to complement on-screen and think-aloud activities.

### *Forced Choice*

Following the CWT, participants were asked to choose between the PDF viewer and A<sup>+</sup>CHIS regarding performance goals of system use. This evaluation method is known as *forced choice*. The following nine performance goals were evaluated: ‘Would you rather use Adobe Acrobat Reader or the A<sup>+</sup>CHIS to (a) get an overview of the domain, (b) develop a general understanding on T2DM, (c) search for specific keywords, (d) capture the main content, (e) search for specific images, (f) get an overview of the most informative images, (g) efficiently navigate through different topics of the content, (h) get answers to questions you might have in mind, and finally, (i) trace past searches’. We decided to choose a *forced choice* rather than a questionnaire format with Likert scales for two main reasons: first, as a formative and explorative study we aimed for an explicit comparison between A<sup>+</sup>CHIS and a challenging benchmark with regards to the nine above mentioned performance goals. Deciding between two alternatives might be

Task	Parallel Version 1	Parallel Version 2
WC1	Which contents/subject areas do you think are included in chapter 5?	Which contents/subject areas do you think are included in chapter 3?
WC2	What is the waist size that creates a greatly increased risk for men?	What daily amounts of beer/alcohol are just acceptable for men?
WC3	How does type I diabetes mellitus develop?	How does type II diabetes mellitus develop?
WC4	What contents/terms have you searched for so far?	What contents/terms have you searched for so far?
TiB1	How often does the term 'smoking' appear in chapter 1?	How often does the term 'stress' appear in chapter 1?
TiB2	Is it worth reading beyond chapter 2 if you want to learn exclusively about 'diastole'?	Is it worth reading beyond chapter 4 if you want to learn exclusively about 'medication'?
TiB3	Does chapter 2 give a better insight into the topic of 'care' than chapter 6?	Does chapter 6 give a better insight into the topic of '(diabetic) foot' or 'foot syndrome' than chapter 3?
TiB4	In which chapter would you most likely start if you wanted to find out more about 'blood pressure'?	In which chapter would you most likely start if you wanted to find out more about 'insulin'?
IS1	According to an illustration in chapter 5, is a sugar value of 100 mg/dl alarming or safe?	Which 3 stages of diabetes therapy are shown graphically in chapter 4?
IS2	Which picture in chapter 6 do you think conveys the most relevant information about type II diabetes mellitus?	Which picture in chapter 6 do you think conveys the least relevant information about type II diabetes mellitus?
IS3	Search for a graphic on the topic of 'sugar metabolism'.	Search for a graphic on the topic of 'nutrition pyramid'.

Table 1: The CWT Tasks.

easier in particular for the inexperienced participants of our sample. Second, considering the sample size which seems reasonable for a first formative evaluation but rather small for a summative one, means and standard deviations which could be derived by Likert scales would not allow for inferential statistical methods.

#### *Semi-Structured Interviews*

Lastly, we conducted *semi-structured interviews* with participants to ask open and generic as well as closed and specific questions about A<sup>+</sup>CHIS as well as further inquiries on the comparison between the two systems. The open questions included for example: 'Which system would you rather recommend to a person as a first source of information to get started?' or about the non-linear content exploration in A<sup>+</sup>CHIS such as 'How much does this interactive system encourage you to explore further content?'. More specific and closed (yes-no) questions were if the participants had already seen or used the different components of A<sup>+</sup>CHIS (such as the Word Cloud or Image Slider, etc.)




Figure 3: The evaluation regarding the prior knowledge and usefulness of A<sup>+</sup>CHIS components by the CWT participants.

prior to the session, if they considered them as helpful and if they would like to use them again.

### 4.3 Results and Discussion

In the following section, we outline and discuss the results, starting from a more global evaluation and continuing with more specific results with respect to the components of A<sup>+</sup>CHIS.

#### Global Evaluation

As outlined in the previous section, in the course of the semi-structured interviews, participants were asked closed questions if they had already seen or used a (i) Word Cloud, (ii) Topic Bar, (iii) Tile Bar, or (iv) Image Slider prior to the session. In addition, they were asked if the components were considered as helpful and if they would like to use them again. A ‘yes-answer’ has been coded as ‘1’, a ‘no-answer’ as ‘0’, and indifference as ‘0.5’. Figure 3 shows the results of this global evaluation: while Word Cloud and Image Slider were well known by participants, as around 75% had seen it before, no participant had come across a Tilebar and only one has come across a Topicbar prior to the session. The Image Slider had been used most often before (ca. 63%). Interestingly, the Word Cloud had rarely been used despite being well known. Several participants noted that they were familiar with a Word Cloud as an overview graphic, but not as an interactive element. Regarding helpfulness and future use, more than half of the participants found the Word Cloud, the Tilebar, and the Image Slider helpful and would consider using them again. However, the Topicbar in its current form was not perceived as very helpful (25%), and only one third of the participants (33.33%) would consider using it again.

#### Evaluation of A<sup>+</sup>CHIS Components

The CWT tasks provided information on how the participants used the different components. We observed mixed results regarding the intuitiveness of searching with the Word Cloud. Some participants were not hesitant to use the Word Cloud and performed efficient searches for keywords and contents, as one participant explained: *I simply*

*looked at the largest terms'* (P02). Others preferred the use of more familiar features, such as the Table of Contents or the Searchbar. However, the Word Cloud seemed to be a tool that was easy to learn. We observed that participants often started to use the Word Cloud when other features were not perceived as helpful due to the nature of the task or because the traditional search was simply inefficient.

Similarly, we observed that the Tilebar component was mostly intuitive and easy to learn. Several participants were completely unfamiliar with the Tilebar but most were able to quickly understand its function and extract information from it, as this quote demonstrates: *'There is a box [tile] in chapter 2 and then no more boxes in the chapters after that, so it [the searched term] is not mentioned anymore'* (P06). Overall, the use of Tilebar allowed the participants to efficiently search for the desired information. With regard to the design, the participants identified a need for optimization; for example, more contrasting colors for the tiles and a clearer labeling of the chapters within the Tilebar would promote more intuitive use.

Finally, the Image Slider was perceived as rather intuitive and efficient for finding images because a number of participants already knew it from other applications and the operation was thus familiar. However, using the search bar was often preferred over browsing the Image Slider, particularly because the latter was time-consuming in chapters with many images, as this participant explained: *'I don't want to click through all of them; they are unnecessary and very small. They don't contain any information; they are just photos without text'* (P07). This quote also highlights the desired reduction of unnecessary images as well as the wish for more context for those images that were deemed not self-explanatory, as expressed by several participants.

### *Performance Goals*

Besides describing the sample as a whole, as part of our explorative research, we compared the results of several subgroups: (i) female vs. male participants, and - based on a Median-split - (ii) 'younger' vs. 'older' participants, 'higher' vs. 'lower' levels of self-assessed (iii) prior knowledge of T2DM, (iv) computer and software skills, and (v) experiences with visualizations. Only the comparison between the two age-groups shows some differences at face level. All other pairs of 'sub-groups' are rather similar in their decisions with regards to the *forced choice* items and thus, we do not report the according results. Figure 4 illustrates the findings of the forced choices participants had to make between the PDF viewer and A<sup>+</sup>CHIS regarding defined performance goals. If participants chose A<sup>+</sup>CHIS as their information source of choice for a performance goal, it received a value of '+1'. Conversely, '-1' was the value for their choice of the PDF viewer and '0' for an indiscriminate choice. Thus, the ordinate ranges from '-12' (all participants chose the PDF viewer) to '12' (all participants chose A<sup>+</sup>CHIS). In Figure 4, the values for the whole sample are reported with green circles, while grey circles represent only the older participants ( $n = 6, \geq 38$  yrs.) and orange circles only the younger ones ( $n = 6, < 38$  yrs.).

The results show the potential of A<sup>+</sup>CHIS with regard to the fulfillment of the performance goals. The values for all nine performance goals are either close to the horizontal 'middle-line' (indicating that participants are equally inclined towards A<sup>+</sup>CHIS and PDF viewer and/or are indifferent) or clearly above, such as for *efficiently navigating* through different topics of the content, getting an *overview* of the most *important images*, or *tracing searches*. We could also observe differences between the two age groups; however, this should not be over-interpreted due to the small sample size but should be given consideration in future research. Overall, A<sup>+</sup>CHIS is evaluated as a good tool for

information searches compared to the PDF viewer, despite participants' familiarity with the latter.

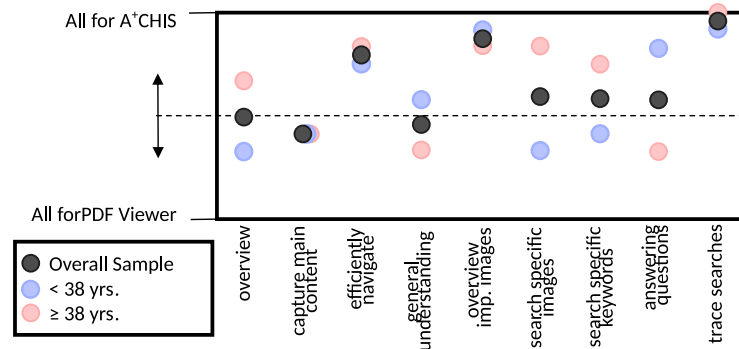


Figure 4: The results of the forced choice evaluation between  $A^+$ CHIS and the PDF viewer w.r.t. the defined performance goals. The performance goals in information processing are sorted along a continuum from more abstract (left) to more specific goals (right).

#### Non-Linear Content Exploration

One area of particular interest was to find out how non-linear content exploration was received by participants. Think aloud during the CWT tasks and the subsequent interviews showed that the system arouses curiosity and motivates further exploration of system features and contents. The participants perceived the search as enjoyable and found the aesthetic design with colors appealing. This applies in particular to the Word Cloud, which manages to cultivate curiosity about further topics, as one participant explained: *'[...] I can imagine that if there is such a Word Cloud I would at least take a look at what topics there are, whether I missed something that would interest me. And I would rather do that than browse through the brochure [...]' (P10)*. It is an added value compared to a non-interactive system, that the Word Cloud fosters engagement with the content further via making interesting and frequent terms more visible.

However, some participants felt that they could not explore the content effectively because they were unsure about how to use the system due to its novelty. In particular, the lack of a familiar linear structure was found confusing and made it difficult for some participants to keep track of the content. As one participant noted: *'[...] I'd like to know what's in it. I don't know what to expect, what the tool offers me. It doesn't give an impression of what it actually is now. At least not a quick one. I don't have that much patience for it [...]' (P07)*. It seems that the open structure overwhelmed some participants, leading to a sense of inadequacy to complete the information search. More support through assistance in the  $A^+$ CHIS system could help users overcome these initial uncertainties so that they can capitalize on the strengths mentioned above and utilize the advantages to the fullest.

## 5 Interaction Analysis

Our data collection efforts, including the on-screen and audio recordings as well as the application of the think-aloud method (Section 4.2), yielded valuable insights and enabled us to evaluate the interaction data of the 12 participants during the 11 CWT tasks in more detail. Thus, the analysis of these interaction data is a secondary analysis of the data obtained during the CWTs as part of the formative evaluation described in the previous section.

These interactions, such as clicks, scrolls, and key-presses, which are performed in order to derive a new insight, are referred to as *(insight) provenance* [Gotz & Zhou 2009] and are a vital cue for the analysis for cognitive processes (Section 5.1). After describing the steps involved in capturing our provenance data such as tools and components used by the participants and the processes involved (Section 5.2), the obtained records were analysed on a quantitative basis (Section 5.3). Lastly, we will show how an interactive visual analysis of such provenance data can be enabled through custom-made visual analytics systems (Section 5.4).

### 5.1 Background

The underlying working hypothesis of the analysis of interaction data was inspired by behavioral mapping, by process models in the field of information seeking and retrieval (e.g., Joseph et al. [Joseph et al. 2013]), as well by the work of Pohl et al. [Pohl et al. 2016] who applied a lag-sequential analysis (e.g., Bakeman and Gottman [Bakeman & Gottman 1997]) to investigate interactions, sequences of interactions and users' activities and processes when engaging with a visualization system. Behavioral mapping is a well-established research method where the paths, movements, and activities within a physical space of participants (when carrying out certain tasks) are recorded and transferred onto a map for further analysis. As an example, Shepley [Shepley 2002] investigated the ways of staff members and their time spent on walking from activity to activity at a neonatal intensive care unit. One of the goals of behavioral mapping is to re-arrange the physical space and workstations to avoid unnecessary paths and to make the workflow more efficient. The underlying principles, questions, and metrics have been transferred to the virtual space of the A<sup>+</sup>CHIS platform and the CWT tasks, having in mind exploratory research questions such as: *What are the processes and interactions of the participants? Are there efficient and inefficient participants with regards to specific tasks, what are their characteristics and how could we support this sub-group of users? Could inefficient series of loops and cycles be avoided by a re-arrangement of the tools and components within the platform, by providing more guidance support, or by highlighting central tools and components which are used by a majority of participants across several tasks?* To answer such questions as these, quantitatively and by means of metrics from behavioral mapping [Ng 2016] and graph theory (e.g. *centrality*), a formal description of the interaction data of each CWT task per participant, the tasks across all participants, as well as the participants across all tasks are required.

The interaction data of one or more participants or for one or more tasks can be graphically represented as a directed, labeled multi-graph, with the tools as vertices and the processes that lead from one vertex  $n$  to a consecutive vertex  $n + 1$  as edge-labels (or vice versa).

## 5.2 Tools and Processes

Before the data processing of the on-screen and audio recordings, a comprehensive list of (cognitive) processes have been defined. Overall, 28 processes have been pre-defined by the three investigators of the formative evaluation study, whereas two additional processes had to be included during the coding process to ensure that all processes are covered by the raw data. Some examples of these processes include *commenting* (when evaluating pictures), *reading* (if text passages have been reproduced verbatim), *interpreting* (concerning pictures and text passages, if participants summarized or evaluated the content in their own words), *pause* (if participants did not do or say anything), but also more basic processes when interacting with the A<sup>+</sup>CHIS platform, such as *scrolling*, *sliding* (through the pictures of the IS), or *hovering* and *click on* (e.g. a certain term in the WC). The tools and exploration subsystems (Section 3), have been further distinguished by considering also the chapter of the brochure [Baumgart et al. 2021] into account; i.e., instead of distinguishing ‘only’ between Image Slider, Word Cloud, etc., it has been further differentiated between the preview (small) and the enlarged Image Slider as well as the Word Cloud, Tilebar, Snippets, etc. for each of the seven chapters of the brochure, resulting in 86 ‘tools’. The coding of the on-screen and audio recordings has been to a large extent carried out individually by the above-mentioned investigators. However, in cases the investigators were not absolutely sure on how to interpret a participants activity and to code it into one of the pre-defined processes, he or she noted the time-stamp and the coding of such ambiguous activities has been done collaboratively by reaching a consensus. The processes have been coded only if they exceeded a duration of around 1 s. For the computation of several metrics and further analysis, the individual interaction data for each of the tasks has been represented as sequence of  $\langle tool_{src}, process, tool_{tar} \rangle$  triples.

## 5.3 Results and Discussion

Overall, i.e., across all 12 participants and 11 CWT tasks, 1,870 processes (=edges) have been observed, whereas 21 out of the 30 processes were applied at least once by the participants. The three most frequently applied processes are *scrolling* ( $n = 268$ ), *sliding* ( $n = 240$ ) and *click on* ( $n = 235$ ). On average, a process took 4.18 s ( $Mdn = 2$  s,  $SD = 5.03$  s), ranging from 1 to 64 s. The tasks (Table 1) where the most processes have been applied by all participants are *IS2* ( $n = 337$ ), *WC2* ( $n = 267$ ) and *IS3* ( $n = 242$ ); whereas the tasks with the least amount of processes are *WC4* ( $n = 65$ ), *TiB1* ( $n = 77$ ) and *WC1* ( $n = 94$ ).

With regards to the ‘tools’, 75 out of the 86 differentiated tools have been applied at least once by the participants, whereas the three most frequently applied tools are the *IS<sub>l</sub>* (the Image Slider in enlarged form) *for chapter 6* ( $n = 308$ ), the *Snps for chapter 1* ( $n = 124$ ) and the *Searchbar* ( $n = 103$ ). To put these numbers in context, the 75 tools which have been applied at least once, were visited 2,002 times across all participants and CWT tasks. The three most applied tools, i.e., those with the highest *centrality*, cover around a quarter of all incoming and outgoing edges.

The triple sequence  $T$  allows to easily evaluate further metrics, such as the number of

1. *loops*, i.e., where the  $tool_{tar}$  is identical to the  $tool_{src}$  ( $|\{ \langle s, \cdot, t \rangle \in T : s = t \}|$ ),
2. *multiple edges*, i.e., where two or more triples have the same  $tool_{src}$  as well the same  $tool_{tar}$ , regardless of the specific *process* that leads from one to the other ( $|\{ \langle s, \cdot, t \rangle \in T : (\exists \langle s', \cdot, t' \rangle \in T \setminus \{ \langle s, \cdot, t \rangle \}) [s = s' \wedge t = t'] \}|$ ), and

3. *identical triples*, i.e., where two or more triples contain the same *processes* as well as  $tool_{src}$  and  $tool_{tar}$  ( $|\{\mathbf{t} \in T : (\exists \mathbf{t}' \in T \setminus \{\mathbf{t}\})[\mathbf{t} = \mathbf{t}']\}|$ ).

Overall, 1, 870 triples have been identified. 1, 030 of them are *loops*, comprising 59 (of the 75) tools. The three tools with the most *loops* are the *IS<sub>I</sub> for chapter 6* ( $n = 288$ ), the *Snps for chapter 1* ( $n = 91$ ) and the *Snps for chapter 3*. These three cases are identical to the three most prominent *multiple edges*. 1, 479 triples represent *multiple edges* with 197 unique tool combinations. Finally, 1, 177 triples represent *identical triples*, with 171 distinct instances. The three most prominent ones are subsets of the above mentioned *loops* and *multiple edges* concerning the tool *IS<sub>I</sub> for chapter 6*, with the processes *sliding* ( $n = 170$ ), *commenting* ( $n = 53$ ) and *viewing* ( $n = 41$ ).

The fact that certain chapters of the brochure occur several times within the top-three of the applied tools, the *loops*, *multiple edges* and *identical triples*, is of course caused by the concrete CWT tasks (e.g., the tasks TiB3 and IS2 which specifically ask for a comparison between the content in chapter 6 and other chapters; see Table 1).

The transitions between tools (agnostic with regards to chapters) are shown in the matrix in Figure 5. The *loops* correspond to the main diagonal, while all cells with a value  $> 1$  contain *multiple edges*. Please note that several transitions from one tool to another are technically not possible; in particular, several transitions from or to the *IS<sub>I</sub>*. When distinguishing only between the more broadly defined 10 tools as in Figure 5, out of the 1, 870 triples, 1, 259 are *loops*, whereas all tools are affected, 1, 802 triples represent *multiple edges*, with 58 distinct tool combinations, and 1, 616 represent *identical triples*, with 221 unique instances. The three most observed *identical triples* are  $\langle IS_1, sliding, IS_1 \rangle$  ( $n = 221$ ),  $\langle WC, scrolling, WC \rangle$  ( $n = 77$ ), and  $\langle WC, scanning, WC \rangle$  ( $n = 74$ ).

The often observed *loops*, *multiple edges*, and *identical triples* with regards to the *IS<sub>I</sub>*, reveal some potential for improvement to make the information search for users more efficient. One potential improvement approach has been suggested by a participant in the course of the formative evaluation study (*P01: An Image Tile Display* that allows for more efficient scanning through the images of a certain chapter at once).

## 5.4 Provenance Visualization

While a ‘global’ evaluation of the provenance – e.g., obtaining quantitative measures for *loops* or usage of individual components – can be conducted based on the raw interaction transcripts, a more in-depth analysis requires dedicated visualizations. Those would allow the analysis of interactions on a per-user and/or per-task basis and could answer additional questions, such as ‘Are the different groups of users observable?’ (e.g., depth-first search vs. breath-first search exploration process), ‘Are there any outliers?’ (users whose exploration process differs significantly from all others), etc. The data’s underlying graph structure invites the application of different established visualizations. We implement two customized visual analytics tools, a graph as well as a matrix layout, which allow us to investigate different orthogonal aspects of the interaction data.

### *Provenance Graph*

In a *Provenance Graph*, we visualize a user’s alteration and switching between different processes; i.e., in a directed weighted graph, we show how often a user switched between different processes (Figure 6, bottom). With this type of visualization, it is possible to, for instance, spot if a user goes back and forth between two different processes or if they go through the same sequence of processes over and over (cycles). We indicate the time




Figure 5: The adjacency matrix of transitions between the ‘high-level’ tools. The entries in the main diagonal corresponds to our notion of loops, while multiple edges populate the remaining cells. The cells reflecting transitions which are technically not possible are intentionally left blank.

spent on specific processes through the size of the respective node and, conversely, the number of transitions between processes through the thickness of the respective links.

#### Provenance Matrix

Besides the graph representation, we also visualize the provenance information in a matrix layout where the sorted rows (Startscreen/DL > ToC > WC/HWC/IS > ToB/TiB > Snps/FullT) represent the high-level tools at different levels of visual granularity and abstraction (overview to closeup ~ top to bottom). The columns reflect the different processes, sorted by their type (from basal/technical processes such as ‘scrolling’ to cognitive/psychological processes such as ‘interpreting’). The matrix cells exhibit a color-coding, indicating the overall time a user spent on a tool-process pair. Additionally, we show the sequence of the exploration with arrows spanning consecutively visited cells. As the display of ‘all’ arrows would overload the visualization, we display only the most recent transitions, with the recentness modelled by an alpha-drop-off. Hovering over a cell triggers a tooltip which lists all the interaction triples responsible for said cell; i.e., as opposed to the provenance graph, the matrix is able to reveal correlations between processes and components, i.e., it shows (i) at which levels of visual abstraction a user predominantly operates, (ii) which processes they carry out at which level, (iii) which processes they carry out using which component, and (iv) whether they exhibit a rather vertical (~ depth-first search) or horizontal (~ breath-first search) exploration behavior.

### *Use Case Examples*

Finally, we investigate the effectiveness of the proposed provenance visualizations for the visual analysis of tasks conducted by different users. To this end, we take a look at one exemplar task (WC3, Version 1: ‘How does type I diabetes mellitus develop?’) which should both, encourage an ‘open’ exploration process, and result in comparable interactions amongst users. For illustrative purposes we compare the respective provenance visualizations of two participants (P01 and P12) who showed vastly different interactions in this particular task. It took P01 188 s to successfully complete the task, while P12 required only 63 s. Figure 6 shows the respective provenance graph and matrix, illustrating the interactions captured while working on said task. Even at a first glance it is obvious that P01 underwent a much more laborious exploration as P12 which is in line with our impression during the supervised evaluation.

On a closer look, we can see P12 used only 5 distinct processes without much back-and-forth between any of them. An inspection of the respective arrows in the provenance matrix also reveals that this participant generally moved from basal processes at a high level of abstraction (top left corner of the matrix) to rather cognitive processes at detail level (bottom right corner of the matrix). This is an expected exploration pattern of a competent information seeker.

P01, on the other hand, required not just more time overall, but had a lot more back-and-forth between different technical and cognitive processes, i.e., ‘Scrolling’  $\rightleftharpoons$  ‘Scanning’ and ‘Reading’  $\rightleftharpoons$  ‘Interpreting’. This interaction pattern is also confirmed by the provenance matrix which further reveals that P01 has several movements from low- to high abstraction levels such as from FullT to WC, or from WC to ToC. We take these bottom-to-top patterns as a cue for an individual who ran into an exploratory wrong track, hence they had to backtrack to a higher level in order to find the correct path to their desired information.

The next research focus will be on the automatic analysis and the clustering of users based on their interaction patterns, as this information can be leveraged to propose adequate visualization to them [Gotz & Wen 2009].

## **6 Overall Discussion**

Our exploration system allows users to navigate through documents and adapt the visual representation and level of detail by using well-known visual analysis techniques such as word clouds, topic models, tile bars, and keyword search. The system provides the users with a two-fold document exploration: a traditional linear and non-linear document navigation by switching between content and detail. Furthermore, it allows users to follow both the edited content of a given document (supervised structure) as well as an automatically computed topic models (unsupervised structure). To the best of our knowledge, there are few empirical studies on the cognitive and motivational aspects of using document visualizations (e.g., tile bars and word clouds, with those of linear document readers). Our evaluation is a first confirmation that our approach could foster interest and heighten curiosity by using a distant-reading approach for exploring the content of interest more efficiently. Further key findings are that the participants in general enjoyed the non-linear content exploration, even if some participants would have needed more support functions at least in beginning of use. There were mixed results regarding intuitiveness of the Word Cloud whereas the Tilebar and the Image Slider have been evaluated as being intuitive; which is also reflected in their agreement to the question if they would use these components again.

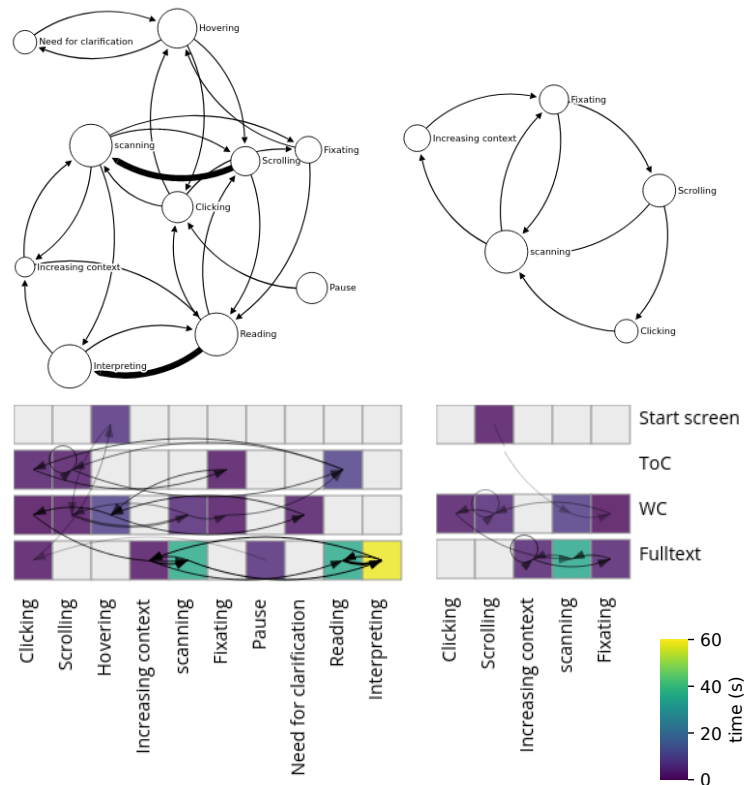


Figure 6: Provenance visualizations of two users, P01 (left) and P12 (right), working on task WC3. Their respective provenance graphs (top) reveal the alternations between processes, while the provenance matrices (bottom) clearly show the alternations between levels of visual granularity.

Our study showed that users did not make great use of the topic model structure. This may be partly due to the unfamiliar representation of topic models. Recently, some studies have investigated the impact of word clouds for topic understanding [Smith et al. 2017] and keyword summaries [Felix et al. 2018]. Word clouds, as it transpired, are particularly useful for quickly identifying the most common and frequent terms, while disadvantages may arise in decoding numeric values from font sizes for larger sets of keywords. An alternative visualization could be a simple word list with a frequency encoding (e.g., font size, bars) that represent each topic. Figure 7 shows such an alternative representation (2) to improve topic understanding where the keywords for each topic are displayed among each other. By using such a layout, overlapping term composition may be included to increase topic understanding and cause confusion as with the previous layout (1). Advanced topic model visualization will be considered in future work.

One key element of our system is its “user tracking” during document exploration. In particular, we track which keywords have been explored, with which visual component, and for how long. This is considered *important information provenance data* which, in our system, is used for provenance visualizations (e.g., history word cloud, interaction matrix



Figure 7: Adaptive word cloud layouts: for improving topic understanding the arrangement can be switched from Wordle layout (1) to an alternative list representation (2).

and graph). The latter is an important functionality for content recommendation, and forms the basis for the mitigation of cognitive biases and potentially harmful and wrong pre-conceptions. The provenance visualizations may reveal emerging difficulties of users during information seeking tasks by showing outstanding patterns of user behavior, e.g., horizontal, vertical or loop-like patterns. Our evaluation is a first step in this direction. To investigate user behavior during the CWT tasks, we integrated the CWT tasks in our provenance visualizations.

## 6.1 Future Work

The advantages of aggregated document representations warrant further examination in future research; in particular, the specific benefits that can be derived from using aggregated document representations. Furthermore, our future work will involve exploring potential misunderstandings that may arise from the highly aggregated nature of certain content presentations.

As evident from Section 5, a vital cue which we plan to leverage for adaptive visualizations are *user interactions*. Research has demonstrated that user interactions can effectively be utilized for recommending particular types of visualizations [Gotz & Wen

2009] and even help mitigate cognitive biases [Gotz et al. 2016]. An investigation of the user interactions obtained from the CWT (Section 5.3) with customized visual analysis tools (Section 5.4) revealed that meaningful interaction patterns, reflecting different types of users, can be observed. We assume that users can be clustered into cohesive groups using said patterns, which, in turn, reflect a user group's need for specific types of visualizations. Therefore, an open challenge is the robust and continuous tracking and classification of these interactions. While this was done manually for the 12 participants of the CWT, a purely automatic solution is necessary for the long-term. Our prototype already comprises a tracking of the components a user is interacting with. To this end, we leverage the mouse pointer position together with the established assumption that a user's cursor movements are correlated with their gaze [Buscher et al. 2008; Reichle et al. 2006]. The other aspect of our interaction logs – the process a user is occupied with (Section 5.2) – is significantly harder to track automatically. Even though purely technical processes such as 'scrolling' or 'clicking' are trivial to track, the capturing of cognitive processes such as 'reading' or 'interpreting' pose a non-trivial challenge. Yet, even such can be determined using low-level mouse interactions together with well-defined heuristics [Huang et al. 2012; Kirsh 2022].

Additionally, we have planned further user studies. These will focus on the planned enhancements of the system, such as the representation of behavioral patterns, which can support users in reflecting on their information-seeking behavior and, thus, potentially contribute to the detection and prevention of biases, especially confirmation bias. This will include investigating how different ways of presenting the interactions can promote unbiased information-seeking by also taking into account differences between users, such as different visualization preferences, different states of knowledge or, as already mentioned above, age groups. Moreover, the need for support, as found in our current study for some users due to the novelty of the system, will be addressed.

In future work, we also intend to research automatic recommendation and develop adaptation methods based on the current system. Such a further developed system version will also be subjected to an extensive case study in order to evaluate the system in detail.

## 6.2 A Model for Adaptation of Content and Presentation

Our initial motivation was to provide an interactive and adaptive CHIS. These systems should adapt the content and presentation to the users' information needs and preferences. Our document exploration design has a number of variables which could be controlled and adjusted by the system to adapt itself to the users' information need and preferences, and recommend views and content. To that end, we define a model based on the following dimensions along which automatic adaptation can be performed. In future work, we will focus on the prediction of dimensions to use and how to set them for specific users.

### *Content Navigation*

Here, *what to show* when the user is doing a mouse-over on a term of the Word Cloud is to be decided. The principal options include (a) go to the best matching full text position (full detail), (b) show the text snippets of multiple matches (an intermediate detail level), or (c) show the tile bar (lowest level of detail) from which the user may pick a finding location. The system could determine the answer based on the previous selections made by the user, presuming the user has a stable preference. On the other hand, the system could track which background knowledge is already available, and show the higher levels of details for topics which are not well-known.

### *Document Structure*

Our A<sup>+</sup>CHIS design shows the section Word Clouds together with the (a) section headings or (b) topic models computed for each section. These represent both supervised and unsupervised content structures. When the user requests a section, the system might adapt to show either one of these. To this end, the system might predict whether the user prefers the traditional, edited document structure given by the headings, or the computed content structure from a topic model. The latter provides an opportunity to compute the topic models such as to represent the users' interest and background information.

### *Configuration of the History Cloud*

The system may present to the user either a History Cloud of (a) what has already been explored, or (b) what has not been explored yet. The system could predict if the user would want to deepen their knowledge on a particular topic or broaden their horizon into new topics. Depending on the user's intent, the system can choose from which terms to create the history Word Cloud. This might also be a good starting point to mitigate biases once they are detected in the users. It is important to recognize that these are concepts, and thus require specific implementations to monitor, track, and characterize the users' background knowledge, preferences, and interests. In order to explore potential implementations and solutions to realize these concepts, both direct and indirect methods will be examined in our future research work.

## **7 Conclusion**

Currently, a significant number of existing CHIS fall short when it comes to presenting health information in interactive, adaptive, and/or personalized ways. In this regard, interactive document visualization techniques are more conducive to enhancing user engagement and promoting a deeper understanding of complex information, thus providing an amplitude of opportunities for improved support of information-seeking tasks. We presented a novel and innovative document exploration system that allows users to visually explore health documents at various levels of detail and abstraction. We incorporated well-known document visualization techniques, such as Word Clouds and Tilebars, into our design and applied it in the domain of T2DM. We evaluated our implemented system by performing a formative study based on a CWT and compared our approach with linear and close reading. The evaluation results are promising and constructive, and demonstrate that our approach is easy to use and helps in content exploration, further facilitating more interactive content navigation as well as motivating users to engage with our implemented system at different levels of visual granularity and detail. We also presented concepts for possible visual adaptation by collecting user interaction data and visualizing this data in two provenance visualizations to reveal specific information needs as well as reading preferences.

### **Acknowledgements**

This work was funded by the Austrian Science Fund (FWF) as part of the project 'Human-Centered Interactive Adaptive Visual Approaches in High-Quality Health Information' (A<sup>+</sup>CHIS; Grant No. FG 11-B).

## References

- [Alencar et al. 2012] Alencar, A. B., de Oliveira, M. C. F., & Paulovich, F. V. (2012). Seeing beyond reading: A survey on visual text analytics. *Wiley Int. Rev. Data Min. and Knowl. Disc.*, 2(6), 476–492. <https://doi.org/10.1002/widm.1071>
- [Arbeitsgruppe GPGI 2016] Arbeitsgruppe GPGI. (2016). Gute praxis gesundheitsinformation. *Zeitschrift für Evidenz, Fortbildung und Qualität im Gesundheitswesen*, 110-111, 85–92. <https://www.sciencedirect.com/science/article/pii/S1865921715300015>
- [Bakeman & Gottman 1997] Bakeman, R., & Gottman, J. M. (1997). *Observing interaction: An introduction to sequential analysis*. Cambridge university press.
- [Baumgart et al. 2021] Baumgart, J., Viegner, U., & Pohl, C. (2021). Den diabetes im griff: Ein handbuch für patientinnen und patienten mit diabetes mellitus typ 2. *AOK-Bundesverband, Berlin*. <https://www.aok.de/pk/magazin/cms/fileadmin/pk/pdf/patientenhandbuch-diabetes.pdf>
- [Blei et al. 2003] Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3, 993–1022.
- [British Diabetic Association 2022] British Diabetic Association. (2022). Diabetes UK - Know diabetes. Fight diabetes. <https://www.diabetes.org.uk>
- [Bundesministerium für Gesundheit 2022] Bundesministerium für Gesundheit. (2022). Verlässliche Informationen für Ihre Gesundheit. <https://gesund.bund.de/>
- [Bunge et al. 2010] Bunge, M., Muhlhauser, I., & Steckelberg, A. (2010). What constitutes evidence-based patient information? Overview of discussed criteria. *Patient Educ Couns*, 78(3), 316–28. <https://doi.org/10.1016/j.pec.2009.10.029>
- [Buscher et al. 2008] Buscher, G., Dengel, A., & van Elst, L. (2008). Eye movements as implicit relevance feedback. *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, 2991–2996. <https://doi.org/10.1145/1358628.1358796>
- [Cao et al. 2011] Cao, N., Gotz, D., Sun, J., Lin, Y.-R., & Qu, H. (2011). Solarmap: Multifaceted visual analytics for topic exploration. *2011 IEEE 11th International Conference on Data Mining*, 101–110. <https://doi.org/10.1109/ICDM.2011.135>
- [Cao et al. 2010] Cao, N., Sun, J., Lin, Y.-R., Gotz, D., Liu, S., & Qu, H. (2010). Facetatlas: Multifaceted visualization for rich text corpora. *IEEE Transactions on Visualization and Computer Graphics*, 16(6), 1172–1181. <https://doi.org/10.1109/TVCG.2010.154>
- [Felix et al. 2018] Felix, C., Franconeri, S., & Bertini, E. (2018). Taking word clouds apart: An empirical investigation of the design space for keyword summaries. *IEEE Transactions on Visualization and Computer Graphics*, 24(1), 657–666. <https://doi.org/10.1109/TVCG.2017.2746018>
- [Firat et al. 2022] Firat, E. E., Joshi, A., & Laramee, R. S. (2022). Interactive visualization literacy: The state-of-the-art. *Information Visualization*, 21(3), 285–310. <https://doi.org/10.1177/14738716221081831>
- [Görg et al. 2013] Görg, C., Liu, Z., Kihm, J., Choo, J., Park, H., & Stasko, J. (2013). Combining computational analyses and interactive visualization for document exploration and sensemaking in jigsaw. *IEEE Transactions on Visualization and Computer Graphics*, 19(10), 1646–1663. <https://doi.org/10.1109/TVCG.2012.324>
- [Gotz et al. 2016] Gotz, D., Sun, S., & Cao, N. (2016). Adaptive Contextualization: Combating Bias During High-Dimensional Visualization and Data Selection. *Proceedings of the 21st International Conference on Intelligent User Interfaces*, 85–95. <https://doi.org/10.1145/2856767.2856779>

- [Gotz & Wen 2009] Gotz, D., & Wen, Z. (2009). Behavior-driven visualization recommendation. *Proceedings of the 14th international conference on Intelligent user interfaces*, 315–324. <https://doi.org/10.1145/1502650.1502695>
- [Gotz & Zhou 2009] Gotz, D., & Zhou, M. X. (2009). Characterizing users' visual analytic activity for insight provenance. *Information Visualization*, 8(1), 42–55.
- [Government of Canada 2022] Government of Canada. (2022). Health Canada. <https://www.canada.ca/en/health-canada.html>
- [Gu et al. 2017] Gu, Y., Wang, C., Ma, J., Nemiroff, R. J., Kao, D. L., & Parra, D. (2017). Visualization and recommendation of large image collections toward effective sensemaking. *Information Visualization*, 16(1), 21–47. <https://doi.org/10.1177/1473871616630778>
- [Hearst 1995] Hearst, M. A. (1995). Tilebars: Visualization of term distribution information in full text information access. *Proceedings of the SIGCHI conference on Human factors in computing systems*, 59–66.
- [Heimerl et al. 2014] Heimerl, F., Lohmann, S., Lange, S., & Ertl, T. (2014). Word cloud explorer: Text analytics based on word clouds. *2014 47th Hawaii International Conference on System Sciences*, 1833–1842. <https://doi.org/10.1109/HICSS.2014.231>
- [Heller et al. 2006] Heller, J., Steiner, C., Hockemeyer, C., & Albert, D. (2006). Competence-based knowledge structures for personalised learning. *International Journal on E-learning*, 5(1), 75–88.
- [Hollingsed & Novick 2007] Hollingsed, T., & Novick, D. G. (2007). Usability inspection methods after 15 years of research and practice. *Proceedings of the 25th annual ACM international conference on Design of communication*.
- [Huang et al. 2012] Huang, J., White, R., & Buscher, G. (2012). User see, user point: Gaze and cursor alignment in web search. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1341–1350. <https://doi.org/10.1145/2207676.2208591>
- [Isenberg et al. 2017] Isenberg, P., Heimerl, F., Koch, S., Isenberg, T., Xu, P., Stolper, C. D., Sedlmair, M., Chen, J., Möller, T., & Stasko, J. (2017). Vispubdata.org: A metadata collection about IEEE visualization (vis) publications. *IEEE Trans. Vis. Comput. Graph.*, 23(9), 2199–2206. <https://doi.org/10.1109/TVCG.2016.2615308>
- [Jänicke et al. 2015] Jänicke, S., Franzini, G., Cheema, M. F., & Scheuermann, G. (2015). On close and distant reading in digital humanities: A survey and future challenges. *EuroVis (STARs)*, 83–103.
- [Joseph et al. 2013] Joseph, P., Debowski, S., & Goldschmidt, P. (2013). Models of information search: A comparative analysis. *Information Research*, 18(1).
- [Keim & Oelke 2007] Keim, D. A., & Oelke, D. (2007). Literature fingerprinting: A new method for visual literary analysis. *2007 IEEE Symposium on Visual Analytics Science and Technology*, 115–122.
- [Kim et al. 2011] Kim, K., Ko, S., Elmqvist, N., & Ebert, D. S. (2011). Wordbridge: Using composite tag clouds in node-link diagrams for visualizing content and relations in text corpora. *2011 44th Hawaii International Conference on System Sciences*, 1–8. <https://doi.org/10.1109/HICSS.2011.499>
- [Kirsh 2022] Kirsh, I. (2022). Virtual Finger-Point Reading Behaviors: A Case Study of Mouse Cursor Movements on a Website. *Big Data Research*, 29, 100328. <https://doi.org/10.1016/j.bdr.2022.100328>

- [Kratwohl 2002] Kratwohl, D. R. (2002). A revision of bloom's taxonomy: An overview. *Theory into practice*, 41(4), 212–218.
- [Kucher & Kerren 2015] Kucher, K., & Kerren, A. (2015). Text visualization techniques: Taxonomy, visual survey, and community insights. *2015 IEEE Pacific Visualization Symposium (PacificVis)*, 117–121. <https://doi.org/10.1109/PACIFICVIS.2015.7156366>
- [Luckin 2010] Luckin, R. (2010). *Re-designing learning contexts: Technology-rich, learner-centred ecologies*. Routledge.
- [Moretti 2005] Moretti, F. (2005). *Graphs, maps, trees: Abstract models for a literary history*. Verso.
- [Ng 2016] Ng, C. F. (2016). Behavioral mapping and tracking. *Research methods for environmental psychology*, 29–51.
- [Plaisant et al. 1998] Plaisant, C., Heller, D., Li, J., Shneiderman, B., Mushlin, R., & Karat, J. (1998). Visualizing medical records with lifelines. *Conference Summary on Human Factors in Computing Systems*, 28–29. <https://doi.org/10.1145/286498.286513>
- [Pohl et al. 2016] Pohl, M., Wallner, G., & Kriglstein, S. (2016). Using lag-sequential analysis for understanding interaction sequences in visualizations. *International Journal of Human-Computer Studies*, 96, 54–66.
- [Reichle et al. 2006] Reichle, E. D., Pollatsek, A., & Rayner, K. (2006). E-Z Reader: A cognitive-control, serial-attention model of eye-movement behavior during reading. *Cognitive Systems Research*, 7(1), 4–22. <https://doi.org/10.1016/j.cogsys.2005.07.002>
- [Rind et al. 2013] Rind, A., Wang, T. D., Aigner, W., Miksch, S., Wongsuphasawat, K., Plaisant, C., & Shneiderman, B. (2013). Interactive information visualization to explore and query electronic health records. *Foundations and Trends® in Human-Computer Interaction*, 5(3), 207–298. <https://doi.org/10.1561/11000000039>
- [Schiefele & Raabe 2011] Schiefele, U., & Raabe, A. (2011). Skills-demands compatibility as a determinant of flow experience in an inductive reasoning task. *Psychological Reports*, 109(2), 428–444.
- [Schwarzer 2008] Schwarzer, R. (2008). Modeling health behavior change: How to predict and modify the adoption and maintenance of health behaviors. *Applied psychology*, 57(1).
- [Shao et al. 2023] Shao, L., Lengauer, S., Miri, H., Bedek, M., Kubicek, B., Kupfer, C., Zangl, M., Dienstbier, B., Jeitler, K., Krenn, C., Loder, C., Semlitsch, T., Zipp, C., Albert, D., Siebenhofer, A., & Schreck, T. (2023). Visual document exploration with adaptive level of detail – design, implementation and evaluation in the health information domain. *VISIGRAPP (3: IVAPP)*.
- [Shepley 2002] Shepley, M. M. (2002). Predesign and postoccupancy analysis of staff behavior in a neonatal intensive care unit. *Children's Health Care*, 31(3), 237–253.
- [Smith et al. 2017] Smith, A., Lee, T. Y., Poursabzi-Sangdeh, F., Boyd-Graber, J., Elmqvist, N., & Findlater, L. (2017). Evaluating visual representations for topic understanding and their effects on manually generated topic labels. *Transactions of the Association for Computational Linguistics*, 5, 1–16. [https://doi.org/10.1162/tacl\\_a\\_00042](https://doi.org/10.1162/tacl_a_00042)
- [Standl et al. 2019] Standl, E., Khunti, K., Hansen, T. B., & Schnell, O. (2019). The global epidemics of diabetes in the 21st century: Current situation and perspectives. *Eur J Prev Cardiol*, 26(2 suppl), 7–14. <https://doi.org/10.1177/2047487319881021>
- [Steele & Iliinsky 2010] Steele, J., & Iliinsky, N. (2010). *Beautiful visualization: Looking at data through the eyes of experts*. "O'Reilly Media, Inc."

[Strobelt et al. 2009] Strobelt, H., Oelke, D., Rohrdantz, C., Stoffel, A., Keim, D. A., & Deussen, O. (2009). Document cards: A top trumps visualization for documents. *IEEE Trans. Vis. Comput. Graph.*, 15(6), 1145–1152. <https://doi.org/10.1109/TVCG.2009.139>

[Sweller 2005] Sweller, J. (2005). Implications of cognitive load theory for multimedia learning. *The Cambridge handbook of multimedia learning*, 3(2), 19–30.

[Wang & Laramee 2022] Wang, Q., & Laramee, R. S. (2022). EHR STAR: the state-of-the-art in interactive EHR visualization. *Comput. Graph. Forum*, 41(1), 69–105. <https://doi.org/10.1111/cgf.14424>

[Wu et al. 2011] Wu, Y., Provan, T., Wei, F., Liu, S., & Ma, K.-L. (2011). Semantic-preserving word clouds by seam carving. *Computer Graphics Forum*, 30(3), 741–750. <https://doi.org/10.1111/j.1467-8659.2011.01923.x>