Christian Wolfgang Burtscher

Concept and Adaptation of Levels of Detail for a CHIS

BACHELOR'S THESIS
to achieve the university degree of
Bachelor of Science

submitted to
Graz University of Technology

Supervisor
Univ.-Prof. Dipl.-Volksw. Dr.rer.nat. M.Sc. Tobias Schreck
Dipl.-Ing. Dr.techn. B.Sc. B.Sc. Stefan Lengauer
bak. mag. Mariia Tytarenko
Institute of Computer Graphics and Knowledge Visualisation

Graz, September 2024

This is a slightly extended version of the official Bachelor's Thesis submitted at Graz University of Technology on 22.08.2024.

Abstract

Consumer health information systems have become very prevalent over the years and have already shown massive benefits to electronic healthcare. This thesis aims to address the adaptation of levels of detail for visualisation techniques used in visualising large amounts of textual data, specifically medical information in a Consumer Health Information System (CHIS). The focus lies on showing possible approaches on how to properly adapt the visualisations used for such medical data while considering individual levels of visual literacy shown by different users. This thesis aims to provide a new method for allowing users to manually change the level of detail of the utilized visualisations and view components, as well as letting the system adapt them automatically using a rule-based approach. The contents of this work will specifically address and adhere to the project A+CHIS, which is an adaptive CHIS that aims at making visualised medical information interactive while also applying the concept of adaptation to realize the adjustment of the system to various user needs.

Contents

CHAPTER 1

Introduction

The A+CHIS [1] is an adaptive consumer health information system that aims to simplify and ultimately enable the study of different medical documents by visualising them using various, mostly text-based visualisation techniques. These visualisation techniques simultaneously help summarise and highlight the most essential sections of the document while also allowing the user to adapt, change and interact with them in order to receive a more individually tailored experience. This should ultimately allow every user to infer the necessary and important information from a given document without having to read through the entire content.

Understanding a visualisation means going through a multi-level process in which the user tries to extract and extrapolate facts, understand how those facts apply and draw conclusions about the data at hand [2]. In order to adhere to multiple user types as well as different levels of visual literacy represented by these users, various levels of detail (LODs) for the applied visualisations are necessary. These LODs are intended to either simplify a visualisation for a more lightweight and easier approach to understanding the documents' contents or increase the complexity of the visualisation and in turn, also be able to offer more input. This ultimately increases the amount of information a user can possibly gather from the currently visualised document.

As it is hard to define which level of detail (LOD) fits a specific user, a visual gallery, which provides a small preview of the supported LODs, can be displayed. This preview is supposed to help the user decide and quickly get a grasp of what a certain LOD changes for a given visualisation. A gallery such as this one could be called a "level of detail gallery" or LOD-Gallery and its analysis as well as its implementation are two of the primary goals of this thesis.

Furthermore, an approach for automatically changing the LOD of a visualisation using a specific set of rules is being addressed. These rules are based on given observation times of the view components used for visualising the medical documents, as well as observed mouse interactions, specifically mouse clicks, with these visualisations. The automatic change in complexity aims to provide suggestions to further enhance the user's efficiency with the system.

1.1 Need for adaptive LODs in a CHIS

E-Health has become extremely important and beneficial over the years leading up to today [3]. Many people and patients in need of medical assistance or information are able to receive this kind of help from multiple different applications. The A+CHIS tries to further improve and increase the usefulness of services that provide electronic health care by expanding the pool of methods for gathering and receiving medical information. The main motivation of this thesis is to extend the A+CHIS, allowing users to work with the system in additional ways in which they can adapt the complexity of the visualisations used to gather medical information quickly and reliably. For that reason, this work directly impacts the quality and support of this type of application and tries to advance the general usability and variety of opportunities that E-Health can provide.

1.2 Research Questions and Contributions

The work of this thesis focuses on further exploring possibilities in changing and adapting LODs for different types of visualisations. The major contribution and research question is whether a matrix approach for displaying these LODs enhances the visibility and usability of the system and provides a better overview of possible adaptations in order to let users quickly and most importantly correctly decide on which LOD fits them the best. With the help of this approach, users can manually change a given visualisation to a different and individual LOD. This allows for a tailored experience and maximizes the information gathered by all visualisations in use.

Additionally, automatic adaptation of LODs for these visualisations is being researched and worked on. A rule-based system is put into place in order to determine a change in the LOD for a visualisation, contributing to general automatic design approaches for adapting user interfaces. Most of the visualisations used in the context of this work are text-based visualisations, as they represent the core of exploring medical documents regarding the A+CHIS.

1.3 Thesis Structure and Organisation

The major part of the implementation of this work has been executed as part of a project directly relating to this thesis, while the focus of this thesis itself lies in explaining and covering the research questions tied to the implementations.

This thesis will touch on important keywords and scientific terms that are important to understand in the context of this work in Chapter 3 before going over the basis of the implementations for both the manual and automatic adaptation of LODs in Chapter 5.

A simple overview of the applied visualisation techniques in the A+CHIS will be provided in Chapter 4 and put into perspective to show why additional levels of complexity are needed and why a novel approach for adapting these complexities is necessary. In Chapter 7, a conclusion of this thesis and a future outline with possible improvements to this work will be presented.

CHAPTER 2

Related Work

The majority of this work is based on and related to previous research which has already been conducted for the A+CHIS. The A+CHIS is a scientific project which entails many different research areas and has thus accumulated lots of information based on interaction as well as adaptive measurements to improve CHIS. Furthermore, this thesis mainly focuses on manually and automatically adapting LODs for visualisations or view components which are present in the A+CHIS. Related work can also be found in automatic adaptation features from various domains.

The work "Visual document exploration with adaptive level of detail – design, implementation and evaluation in the health information domain." [1] by Shao et al. represents the primary source of context and is used as a baseline for this thesis. In this publication, a precise and detailed overview of the A+CHIS is presented and the overall structure of used visualisations and their interactive features is conceptualized. Other work, such as "Hierarchical Topic Maps for Visual Exploration and Comparison of Documents." [4] by Tytarenko et al. or "Recognizing User Behavior from Interactions for Adaptive Consumer Information Systems." [5] by Lengauer et al. provide additional insight into how the medical data obtained from documents is modelled in order to generate proper interactive visualisations used for exploring this modelled information and how the tracking and observation of user interactions in the A+CHIS work.

Furthermore, the work "Adaptive visualisation of Health Information Based on Cognitive Psychology: Scenarios, Concepts, and Research Opportunities" [6] by Schreck et al. provides additional insight into why adaptive visualisations in a CHIS could improve health literacy while also increasing the effectiveness of the general medical system [6].

Lastly, the work "A Context-Based Adaptive visualisation Environment" [7] by Golemati et al. as well as "Model-based adaptive user interface based on context and user experience evaluation" [8] by Hussain et al. also talk about the necessity of automatic adaption in modern visualisation environments and further endorse the process of automatic adaptation based on user interactions with the system such as mouse movements or pressed keys.

CHAPTER 3

Background and Terminology

The immense benefits made possible by electronic healthcare can be cut short if the system or platform the application operates on is not solid enough. This means that important information can be lost or not conveyed properly if a certain user working with the consumer health information system (CHIS) does not fully understand or comprehend the shown data. This goes against the principle of a CHIS, where the main focus lies on providing medical information to people who might not be versed or knowledgeable in a specific medical field but still show interest in selected topics thereof [9]. Medical documents or data tend to be rather long and difficult to understand, which is why a proper CHIS should be able to present medical information to users who might not be able to work with the complexity and size of such medical datasets. Additionally, a CHIS should be able to extract the most important information and make it available to the user without cutting off relevant details.

An essential factor that usually hinders a proper CHIS in providing beneficial health information, is the methodology used to present medical data to the user. How the information is depicted plays a major role in how well a user or patient understands the contents presented and how quickly and efficiently they can infer necessary information from it. In a CHIS, which visualises medical information by making use of different data graphs or visualisation techniques, a user might not be able to infer all the needed information since the visualisation itself might be too complex and difficult for them to work with.

Not only does the amount of possibly retrievable information a user can infer from a visualisation vary from one visualisation method to another, but also from one user to the next. This simply means that different users might not be able to gather the same amount of information in the same amount of time. It is therefore important to recognize this problem and visualise the information in a way where all kinds of users and user groups are able to infer the information they deem important reliably [10]. These typical and mostly inevitable problems related to consumer health information systems directly motivate this thesis and serve as the basis upon which this work and its implementations have been built.

3.1 Visual Literacy

Visual Literacy is the term used to define a person's ability to properly process and infer information from a given set of data graphics or visualisations [11]. Different people might not be able to gather the same amount of information from a given graph in the same amount of time. While some users are capable of inferring all the key information needed by looking at some visualised output of a dataset, others might need more time or might not be able to comprehend the visualisation at all. Visual Literacy can thus be defined as a group of vision competencies a human being can develop by seeing and at the same time having and integrating other sensory experiences [12] and needs to be taken into account when working on a CHIS which tries to provide medical knowledge to different kinds of users with various levels of visual literacy.

Consequently, it is imperative to acknowledge different levels of visual literacy and try to consider them, especially when making use of text-based visualisations for a CHIS. This ensures that many users are able to work with the system in a satisfactory manner and receive the information they might be looking for without missing essential details or facts by not being able to comprehend a certain type of visualisation.

3.2 Level of Detail

A widely used way of adapting visualisations and data graphics in a way that could consider the individual level of visual literacy a user has is by giving these graphics multiple LODs. The term LOD refers to the capability of providing different representations of an object, which in this case refers to the visualisations of medical data, at different levels of accuracy and complexity [13]. There are many ways of implementing LODs for certain types of visualisations or graphics, featuring different approaches and possibilities, ranging from adapting simple nuances to changing the entire design or general look of a visualisation. Various known methods for implementing and allowing a change in a LOD for a visualisation come with their respective positive as well as negative effects.

Choosing an appropriate method for adapting a LOD is usually quite a big problem, as different use cases call for various techniques for changing the detail of a given visualisation [14]. Since the A+CHIS primarily makes use of text-based visualisations and infographics, the typical methods that can be found for implementing LODs might not work properly or achieve the intended effect. This has to do with the nature of these types of visualisations which, most of the time, lack the depth and dimension to allow for generic lens or zooming approaches which are often made use of in 3-dimensional graphs, images or objects with a certain level of depth.

This results in a LOD selection problem requiring a new way of thinking for allowing changes in looks or even design for the visualisations that might be used in a CHIS. A novel approach for changing such LODs will be proposed and discussed in this thesis.

CHAPTER 4

A+CHIS

The A+CHIS is an adaptive consumer health information system which ties in all the basic elements of a regular CHIS while also making the system and the information which can be made available to patients and users interactive and adaptive. The A+CHIS serves as a project which tries to answer a major research question of whether evidence-based medical information and novel interactive data visualisations can be combined to further enhance the positive benefits to healthcare that usual consumer health information systems already provide [1].

This is achieved by making the visualisations, which are responsible for conveying and summarizing the information present in large medical documents, interactive and adaptive. The interaction with the visualisations itself opens up further exploration possibilities, which allows patients and users to further research a selected topic and extract the exact information they might be looking for. The adaptive aspect, on the other hand, ensures that the visualisations used to display medical information are adjusted to fit the users' needs and capabilities, maximizing their information gain without making the visualisations too complex.

A lot of different services and features are put into place in order to be able to achieve the aforementioned degree of interaction and adaptability. The arguably most important nature of the A+CHIS which is also mainly responsible for allowing a change in the LOD for a visualisation to happen automatically is the tracking and observation of most user interactions [5]. The A+CHIS stores and tracks mouse movements as well as mouse clicks and dwelling times on certain view components of the application in order to gain a better understanding of how a user behaves and reacts to changes in the system. This allows for adapting the entire CHIS to further suit a patient's or user's needs and thus enhance the medical benefit that the A+CHIS provides.

The usage data which is locally obtained by observing user interactions with the system is also used as a baseline for determining a fitting LOD when recommending automatic adaptations. Some images and sections in the upcoming chapters will display and cover information used for testing purposes during the implementation of this work. This test content has been sourced from the AOK book "Den Diabetes im Griff" [15] and refers to diabetes, a very complex disease and highly relevant to the public today [1].

4.1 Framework

The A+CHIS makes use of an Angular frontend paired with a Django backend. Typescript is the primary language that is used to develop and implement frontend features. Different basic JavaScript libraries are used to control the flow of the interface, like RxJS which is used to notify components in order to induce some type of change when events that might have been triggered by the system or the user occur.

RxJS and various services which have already been introduced to the system have proven to be invaluable and provide a strong framework for the LOD approach that this work covers. Through user interaction tracking and a high-level subscription system provided by RxJS, it is possible to adapt the LOD of visualisations and view components on demand and during runtime.

4.2 Visualisation techniques

In the following section, all the visualisation techniques which are employed in the A+CHIS in order to visualise the information extracted from medical documents will be introduced briefly and accompanied by figures which show the initial state of a given visualisation prior to this work. Additionally, a short explanation or overview of why the initial version of such a method might not be suitable for different users with varying levels of visual literacy will be provided indicating the necessity for additional LODs.

4.2.1 Word Cloud

A word cloud is a very common and widely used form of text-based visualisation. It appears as a compact cloudlike structure consisting of terms which are often displayed in different font sizes and at different positions used to indicate a certain ranking of importance between the terms where a higher font size correlates to a higher importance of the term occurring inside the main text [16]. In the A+CHIS, the word cloud also makes use of different colourations for terms that belong to the same topic, which insinuates topic modelling [4].

Figure 4.1: A Word Cloud in the A+CHIS with a high LOD.

Looking closely at this word cloud, it is apparent that an adaptation in the form of changing the LOD for this visualisation can prove quite simple since there are not a lot of dimensions present which could affect the overall information load or design of the visualisation. However, for some patients or users working with the system, the level of complexity shown in Figure 4.1 could already prove to be too difficult since a lot of terms are displayed, which end up making the word cloud appear a bit too cluttered.

Due to the simple nature of this visualisation technique, an adaptation to the LOD is quite straightforward. Simply increasing or decreasing the amount of terms shown, as well as reducing the amount of groups present, directly increases and decreases the level of complexity that this visualisation epitomizes.

4.2.2 Topic Cloud

Topic clouds function and appear quite similar to word clouds, with the key difference being that topic clouds inherently group the terms based on topic and visually place the grouped terms apart to better indicate these clusters. In the case of the A+CHIS, they also show it with different colourations for terms belonging to the same topic. The font size of the terms yet again indicates their frequency of occurrence and importance.

Figure 4.2: A Topic Cloud in the A+CHIS with a high LOD.

Similar to the word cloud, the simplicity of this visualisation technique also allows for a rather simple adaptation in its LOD. Merely increasing or decreasing the amount of terms displayed per topic, or reducing the amount of topics entirely, can directly impact the readability and the overall LOD of the visualised contents.

4.2.3 Tilebar

Word clouds and topic clouds are the primary text-based visualisations used to summarize and convey the contents of a medical document. The tilebar is used to additionally enhance the user's ability to further study a few specific terms which have been visualised using the primary visualisation techniques [4]. This is done by hovering over a term contained in either a word cloud or a topic cloud.

This interaction with the visualisation opens up a tooltip showing more specific data about the occurrence of the term inside all available chapters of the document which is being explored at the moment.

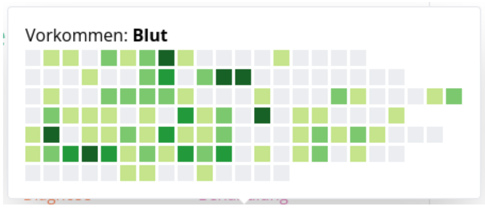

Figure 4.3: A Tilebar in the A+CHIS with a high LOD.

Every single square shown in Figure 4.3 represents the number of occurrences of the term in a specific section of different chapters inside the document [1]. The colouration is once again used to convey this exact information. A darker colour implies a higher amount of occurrences, whereas a lighter colour implies a lower amount of occurrences.

It might seem very difficult for most people to immediately understand and be able to work with this type of visualisation. The various differently coloured squares possess almost no implicit indication of what they intend to show, which can ultimately lead to confusion about the information the user tries to infer. This means that designing different variations of this visualisation is imperative to achieving the desired behaviour of a properly executed adaptation in the LOD. In the case of this tilebar, a lower LOD was achieved by reducing the squares down to actual bars, thus grouping the individual information which every single square had and gathering this information in one collective rectangle. An even lower LOD was achieved by removing the rectangular design of this visualisation altogether and simply displaying one singular sentence which describes the information of occurrence directly.

4.2.4 Text Snippets

Text snippets will appear on the interface after having clicked on a term present in the two primary text-based visualisations. These snippets emerge on the right side of the screen and allow the user to further study selected paragraphs and sections of the document [1].

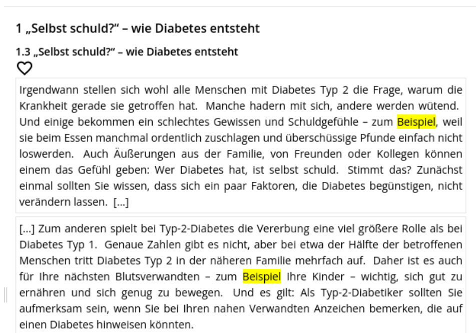

Figure 4.4: Text snippets in the A+CHIS with a high LOD.

This provides the user with a tailored search experience where only the contents which the user was initially interested in will be shown. Since the dimension of complexity to plain text snippets is very comprehensible for most users, a change in a LOD for such snippets can be achieved by simply reducing the number of sentences displayed per snippet. The lowest LOD implicitly relates to only one sentence being shown per snippet.

4.2.5 Thumbnail Slider

The final visualisation technique users can interact with to discover the medical information they deem significant, making up a key view component of the exploration view for displaying information about medical documents, are thumbnail sliders. The thumbnail sliders used in the A+CHIS consist of the infographics and images which might be present in the document at the specific section which is summarized in the corresponding primary text-based visualisation and displayed in a small collection [1].

Figure 4.5: Thumbnail sliders in the A+CHIS with a high LOD.

As with many of the other previously shown visualisations, the type of thumbnail slider used in the A+CHIS does not feature a lot of diverse dimensions, which could make designing a variation of different complexity for an adaptation in its LOD difficult. By simply increasing or decreasing the amount of infographics shown at any given time, the LOD of this visualisation technique can be adapted quite easily. An increase in displayed infographics indicates an increase in its LOD, while the opposite indicates a decrease in complexity and thus a decrease in its LOD as simply more or less information is shown at a given time.

CHAPTER 5

Adaptation

The major focus of this thesis lies in the concept of adapting the LODs for specific view components and visualisations in the A+CHIS. The LOD of a certain visualisation should both be manually and automatically adaptable. Many concepts and ways of changing LODs already exist. Still, none of them are specifically catered to consumer health information systems or text-based visualisations, which makes them hard to integrate and adjust. This is why this work focuses on creating a novel approach to manually adapting such a CHIS and its view components to fit the individual levels of visual literacy of its users as well as making use of the tracking aspect of the A+CHIS to provide fitting recommendations for adjusting LODs automatically.

The purpose of manually adapting the complexity of the view components is to allow the users to decide for themselves whether a preset or recommended LOD is suitable for them or should be adapted further [17]. This way, changing the appearance of all important and necessary visualisations used for exploring documents allows users to find an adequate setting of LODs to comfortably work with the system and be able to maximize the information which could be retrieved by a specific user.

5.1 Manual Adaptation

This way of adapting the complexity of the view components should be feasible and easy to understand and should provide a good overview of what a possible change to the LOD of a visualisation could look like. It is supposed to be comprehensible to all types of users, no matter their level of visual literacy, and therefore facilitate effective decision-making through the mitigation of information overload [18].

The concept of manual adaptation features a matrix approach. This matrix holds all view components and visualisations for which a change in their LOD is supported. This matrix is designed to provide a quick outline of what changes are available to the user and how the visualisation's appearance would alter if a modification in its LOD occurred.

Figure 5.1: The matrix concept in the A+CHIS for manual LOD adaptations.

Figure 5.1 shows the design of the matrix and the way it has been implemented in the A+CHIS. The three columns of this matrix show the available LODs, whereas each visualisation or view component is represented by a row. The blue border which can be seen around an individual item present in each row is representative of the currently set LOD for its view component. This way, it quickly becomes apparent what the active LOD for a specific type of visualisation is and how the other two LODs differ from the one currently in place.

This approach allows for a quick adaptation of a LOD for any given visualisation and provides a preview of what a potentially higher or lower LOD could look like. Additionally, a column-wide selection which sets the LOD of all view components at the same time is possible. This can be considered as a global change and adapts the entire exploration view to the chosen LOD.

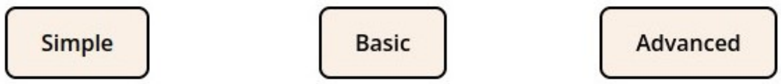

Figure 5.2: The column wide selection options for manual adaptations.

By pressing one of the three buttons, a global adaptation for all visualisations will trigger, changing them to the selected LOD.

This matrix approach ensures that the user or patient working with the CHIS has full control over the complexity of its view components. This ultimately enables the customization of every single visualisation and decouples them from a LOD that another visualisation might have. Given that a user could possess a higher level of visual literacy for a specific visualisation technique than for another, the matrix allows all of the visualisations to be adapted to an individual and suitable complexity and preference setting which is essential for users to understand the contents and properly infer information [19]. This approach thus eliminates the restriction which usual LOD selection methods like global LOD selection or zooming and lens approaches have.

5.2 Automatic Adaptation

The automatic adaptation approach considers the users' interactions with the system as well as the dwelling or observation times gathered for each view component when trying to evaluate possible changes made to the LOD of some visualisations [5]. If the system decides that an adaptation might benefit the user, a recommendation in the form of a modal dialogue will appear. This window will present the adaptations in the LODs which were deemed to fit the user better than the ones currently active, and allow the user to select the changes that they agree with. This method prevents sudden workflow interruptions and does not confuse the user, as the evaluated recommendations are not implemented as soon as the system is finished with its assertion [8]. The user knows what is happening, can take control over the recommendations, and is even able to shut off and turn back on the automatic adaptation recommendations entirely.

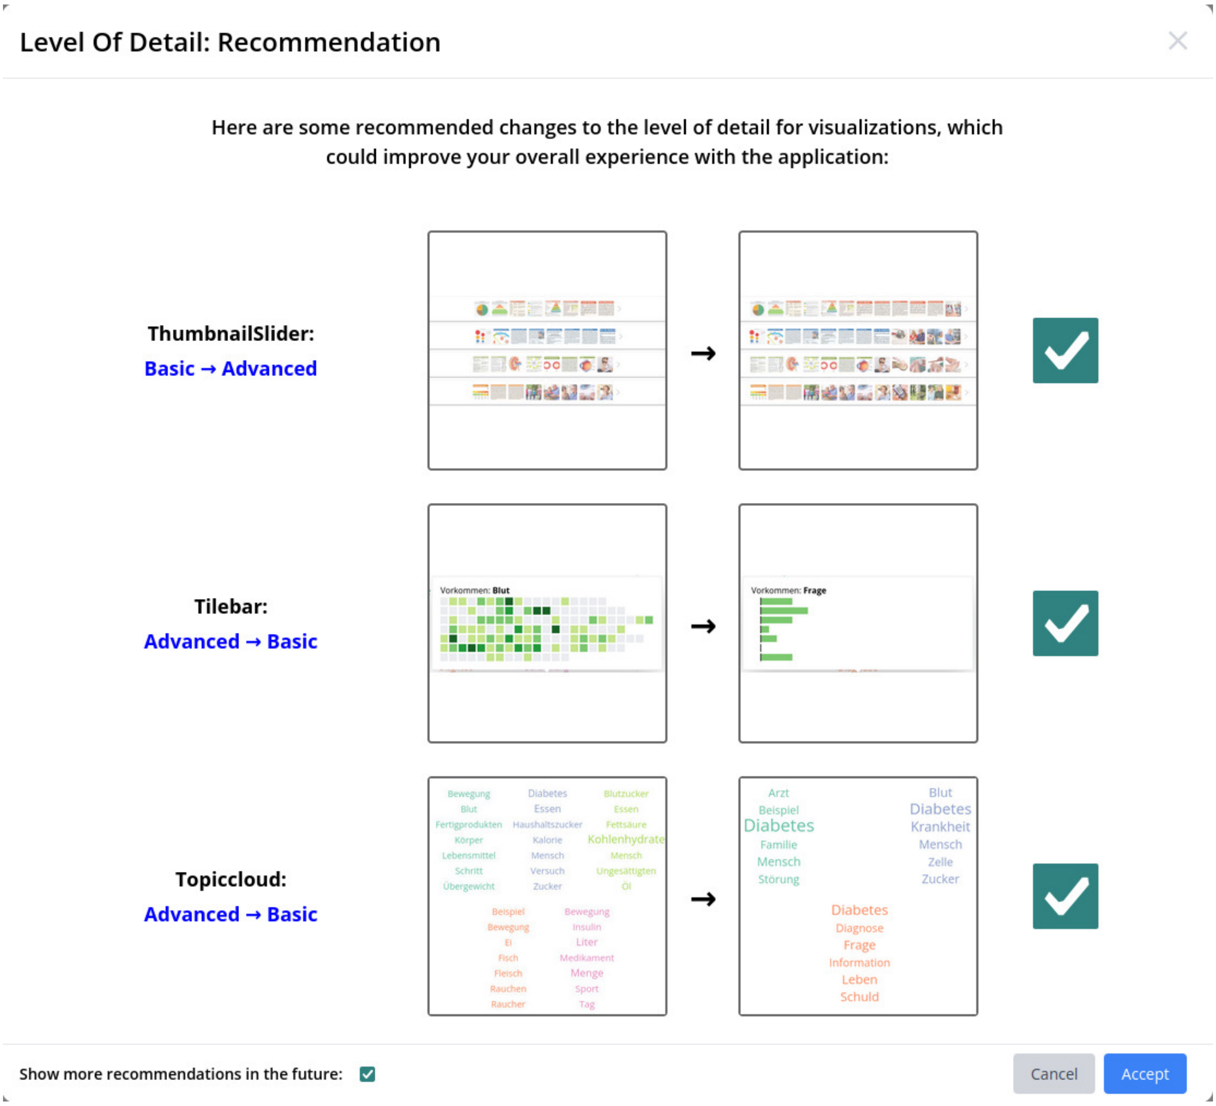

Figure 5.3: The recommendation concept in the A+CHIS for automatic LOD adaptations.

Figure 5.3 shows the implemented concept of automatic adaptation in the A+CHIS. Every entry which is presented on this window consists of the type of view component for which a recommendation is being made, the current state of its LOD and the actual recommended adaptation to which LOD it should adjust based on the systems' perception.

The system enables the selection and deselection of any recommendation it makes. Additionally, at the bottom of the modal dialogue, an option for turning off future automatic recommendations is provided and the proposed adaptations that are currently visible can either be accepted or denied.

This approach allows users with different levels of visual literacy to gain a preview and understanding of what the adaptations would do and how they would change the system's appearance after taking effect. This still allows full control over the adaptations made to the exploration view, while providing valuable input to the user nonetheless.

5.2.1 Accuracy

The actual evaluation of the automatic adaptation makes use of a simple rule-based approach. Since the A+CHIS tracks the user's interactions with the system and observes the amount of time a user works with or looks at a specific type of visualisation [5], a lot of data is available to decide whether a currently active LOD is appropriate for the user.

However, one major problem with making accurate recommendations always remains: The heuristics in a rule-based approach for specifying whether an observation points in a direction that leads to believe that the user is comfortable with a LOD and knows how to work with it are very subjective and could very well be regarded as the opposite for a different user. This implies that there is no definitive way of defining a set of rules which meet all constraints [20] and apply to all types of patients and users possessing various levels of visual literacy.

The defined outcome of the set heuristics might be suitable and related to the behaviour of one user where, for example, a higher amount of mouse interactions, specifically mouse clicks, correlates to more enthusiasm about the LOD of a visualisation and thus increased motivation to use it more often than others, whereas the same behaviour could correlate to another user being too indecisive with the present LOD of the visualisation and thus imply improper interactions resulting in more mouse clicks than usual. This means that recommendations for automatic adaptations could always be erroneous since a user might not directly adhere to the heuristics used for the rule-based system. This is why the rule-based approach of this work generalizes the meaning and indication of high and low dwelling times as well as mouse clicks, recognizing that the applied heuristics might not fit all user scenarios.

5.2.2 Extraction

Before recommendations can be made to the user, the system retrieves all available observed data about the user's interactions and dwelling times [5], extracts only the necessary and valuable information needed and forwards it to the actual evaluation where the heuristics are applied. This automatic adaptation cycle repeats itself continuously in set intervals. Every time a set interval passes and a possible recommendation can be made, all the observed user behaviour from the current interval is retrieved and filtered. Previous observation sessions will not be taken into consideration to keep the latest recommendations clean of any past influence and only consider new information about the user's behaviour toward the system.

Both view durations and mouse clicks are grouped into one final measure per visualisation type. This means that every single instance of a recorded observation will be summed up to a final value which indicates the overall observation duration or dwelling time the user had for a specific type of visualisation. The same applies to the mouse interactions where all recorded instances of mouse clicks on certain view components will be summed up and grouped into a final indicator representing the total amount of mouse clicks per visualisation. Ultimately, these two measures are the representatives used to determine possible changes to LODs and show how long and intensively a user has worked with and looked at a distinct type of visualisation.

5.2.3 Heuristics

The heuristics for the rule-based automatic adaptation proposed in this thesis are rather simple and are split into two parts. Once the total amount of mouse clicks and the overall viewing times of the visualisations have been gathered, the visualisation with the highest view duration will be considered the baseline for the first application of rules. The highest dwelling time, which is represented as an integer in the code, will be divided by three in order to simulate sections, in which the other dwelling times or view durations can be put into. The first set of rules is then applied as follows:

(1) The first section is represented by the higher third of the highest recorded dwelling time. The visualisations with dwelling times that fall into this section are considered easily comprehensible and perceived as simple to the user resulting in these high view durations. Any visualisation which ends up in this section will receive a recommendation for **INCREASING** the current LOD which could result in the user inferring more information without making the visualisation too difficult.

(2) The second section is represented by the middle third of the highest recorded dwelling time. The visualisations with dwelling times that fall into this section are considered appropriate and perceived as suitable for the user. Any visualisation which ends up in this section will receive no recommendation for an adaptation in its LOD.

(3) The third section is represented by the lower third of the highest recorded dwelling time. The visualisations with dwelling times that fall into this section are considered difficult and perceived as incomprehensible to the user resulting in these low view durations. Any visualisation which ends up in this section will receive a recommendation for **DECREASING** the current LOD to make the visualisation easier to work with.

Once the first set of rules has been applied, the mouse interactions are then looked at and evaluated in a very similar manner. This allows the system to consider both measures independently and make adaptations to the visualisations based on their value alone. The visualisation with the highest amount of mouse clicks will be taken as a baseline for the application of these rules and once again divided by three to get three sections in which other visualisations can be put into based on their amount of mouse interactions. The rules are then applied as follows:

(1) The first section is represented by the higher third of the highest recorded amount of mouse clicks. The visualisations with a count of mouse clicks that fall into this section are considered easily comprehensible and perceived as simple to the user explaining the high amount of interactions and willingness to work with this view component. Any visualisation which ends up in this section will receive a recommendation for **INCREASING** the current LOD which could result in the user inferring more information without making the visualisation too difficult.

(2) The second section is represented by the middle third of the highest recorded amount of mouse clicks. The visualisations with a count of mouse clicks that fall into this section are considered appropriate and perceived as suitable for the user. Any visualisation which ends up in this section will receive no recommendation for an adaptation in its LOD.

(3) The third section is represented by the lower third of the highest recorded amount of mouse clicks. The visualisations with a count of mouse clicks that fall into this section are considered difficult and perceived as incomprehensible to the user explaining the low amount of interactions and a lower motivation to work with this visualisation. Any visualisation which ends up in this section will receive a recommendation for **DECREASING** the current LOD to make the visualisation easier to work with.

If a visualisation has already seen a recommendation for either increasing or decreasing its LOD based on dwelling times in the first set of rules, an additional increase or decrease will not be applied if the second set of rules has determined a similar outcome. However, if a visualisation has seen an increase or decrease in its LOD based on the first set of rules, the opposite recommendation can still apply based on the second set of rules. This would result in the user not receiving a recommendation for adapting the LOD of this visualisation after all.

Heuristics		**Trigger**	**Action**
R1) High Interaction & Motivation	Highest recorded value Higher Third	Tracked value is more than or equal to two-thirds of the HRV*.	Recommend **Increase** in LOD
R2) Expected interaction	2/3rd Middle Third	Tracked value is higher than one-third and lower than two-thirds of the HRV*.	No Recommendation (LOD remains the same)
R3) Low interaction & Motivation	1/3rd Lower Third	Tracked value is less than or equal to one-third of the HRV*.	Recommend **Decrease** in LOD

*HRV → Highest recorded value across all observed visualisation types

Figure 5.4: An overview of the heuristics used for automatically adapting LODs.

Figure 5.4 shows a table indicating the three rules which are applied to both the mouse clicks and the view durations. This table is representative of the heuristics for both measurements as they are evaluated independently. The same three rules are applied directly after each other in two consecutive runs with the first being for the viewing times and the second for the number of mouse clicks. The table shows the three sections mentioned above and how the delimiters defined by the visualisation with the highest recorded value across all observed visualisation techniques are used for designating these sections where the other visualisations and their corresponding recorded values can fit in.

A value representing either the number of mouse clicks or the total view duration captured in the observation sessions since the last recommendation was made, triggers the rule for recommending a decrease in the LOD of its corresponding visualisation, if it is less than or equal to one-third of the highest value across all observed visualisations. Similarly, a recommendation for increasing the LOD of its visualisation is triggered if the value is higher than or equal to two-thirds of the highest recorded value. This also indicates that the visualisations with the highest view duration and the highest number of mouse clicks will automatically receive a recommendation for increasing their LOD. Finally, no recommendation for changing the LOD of a value's corresponding visualisation is made if it is both higher than one-third and lower than two-thirds of the highest recorded value.

After having evaluated all possible adaptations to the LOD of the view components, the information is collected and sent to the modal dialogue window which then pops up and ultimately recommends the adaptations to the user. If the user has previously decided not to receive further recommendations, the entire evaluation process will not begin in the first place. This leaves the information gathered from observation sessions in these time frames open for future recommendations and does not exclude them as if the evaluation process had already begun.

CHAPTER 6

Walkthrough and Evaluation

The manual as well as the automatic adaptation ultimately aim at improving the usability of the system by allowing users to change the complexity of the components responsible for visualising the data that can be retrieved from a medical document to fit their information needs [1]. These changes enhance the user's ability to retrieve and infer information by considering their visual literacy. The manual adaptation, which is based on a LOD-Gallery displaying a matrix with all available LODs per view component, allows users to check and decide for themselves which LOD might fit their visual literacy needs the best. This gallery is part of the graphical user interface and can be viewed and used to introduce change to the system at any point in time. The automatic adaptation is designed to recommend adaptations to the complexity of the system by observing user interactions and behaviour with the view components and evaluating them based on a set of rules. These recommendations are featured in a way that does not interrupt the workflow of the user and does not cause confusion by realizing the system's adaptations immediately and frequently [8].

The following walkthroughs provide a short overview of how the adaptations to the LODs of view components can be realized and how they could benefit the user. They are intended to evaluate the advantages of the adaptations and the implementation thereof in order to show how the user is able to better their experience with inferring medical knowledge by properly making use of the visualisations in place and adapting them to their specific needs.

6.1 Manual Adaptation

The LOD-gallery in the A+CHIS is displayed using a modal dialogue which can be opened by simply pressing a button on the graphical user interface. Once the matrix, which provides an overview of all possible adaptations to the LODs for the visualisations as well as the currently active LOD, is visible, the user can utilize the matrix by selecting the design of a visualisation they seem to understand. The selected LOD will automatically receive the blue border mentioned in Section 5.1 to indicate that it is the currently active LOD for its type of visualisation thus showing that the manual adaptation has been realized.

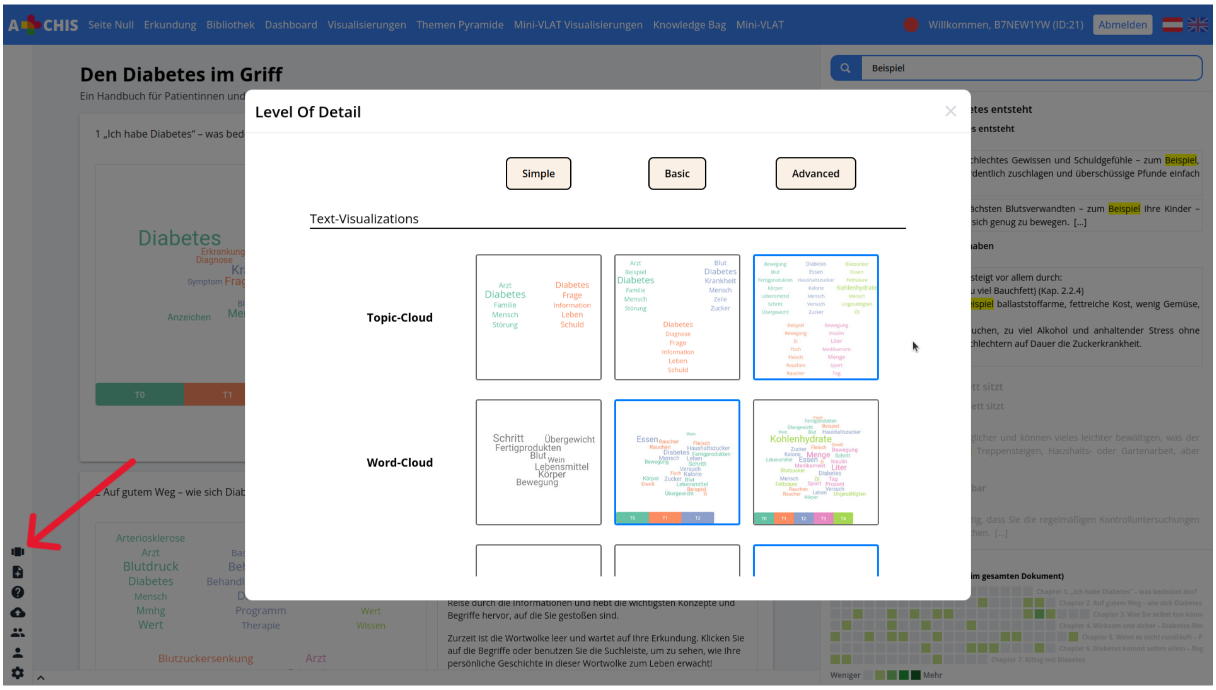

Figure 6.1: The graphical user interface of the A+CHIS with the opened LOD-Gallery.

The interface shown in Figure 6.1 represents the look of the A+CHIS after the modal dialogue containing the matrix for manually adapting the LOD of view components has been opened. The exploration view is still visible in the background. A change in the complexity of its view components can be recognized immediately upon selecting a new LOD. The red arrow in Figure 6.1 indicates the button the user has to press to open the LOD-Gallery. The user can scroll through the matrix and discover all view components with an adaptive nature as well as preview their LOD. This manual adaptation approach supports the user in proper decision-making regarding their own level of visual literacy and attending to their abilities [18], providing a simple but effective solution to introducing change to the system's intricacy.

6.2 Automatic Adaptation

The process of automatically adapting the LOD of certain visualisations starts with the user. By working with the system, utilizing the visualisations to gather information and searching for additional topics after having interacted with the view components, the A+CHIS is able to track and observe the user's behaviour. This data is captured over multiple observation sessions which are finished once the user taps out of the system or changes windows leaving the A+CHIS in the background. After a set time interval of forty-five minutes, the automatic adaptation process starts by extracting the aforementioned observation data and preparing it for evaluation. During this process, the user is still able to work with the system as usual and will not be able to tell that recommendations for automatically adapting the view components are currently being instigated. Once the extraction and evaluation mentioned in Section 5.2 are finished, the modal dialogue shown in Figure 5.3 appears automatically.

Here, the system's recommendations which have been based on previous user interactions are displayed. Every single change in a LOD of a visualisation is represented by a singular row which indicates the proposed change visually and textually. Additionally, each recommended adaptation can either be selected or deselected to ultimately leave the choice of implementing the adaptations to the user.

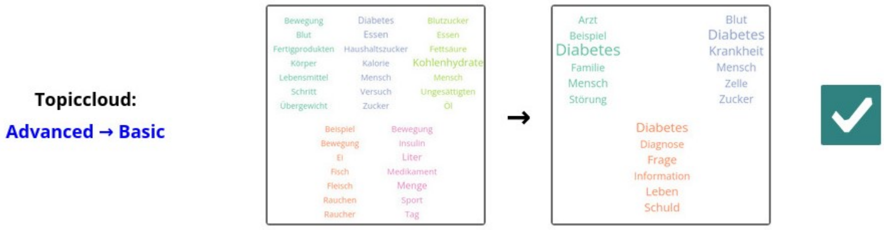

Figure 6.2: A singular automatic adaptation recommendation and its selection.

The method of displaying recommended changes shown in Figure 6.2 allows the user to once again preview what the realization of such an adaptation would do to the affected view component thus supporting the user in ultimately deciding whether a change in complexity should be implemented. This does not take away the control of the user by executing the automatic adaptions without review and setting them in place immediately while leaving the user uninterrupted and focused on the task they are currently pursuing [18].

Figure 6.3: Options for accepting automatic adaptation recommendations.

Finally, after the user has selected all the adaptations they deem plausible, they can still decline all recommendations and leave the system as is or accept the changes for the chosen view components. Furthermore, the user is given a choice of whether future automatic adaptation recommendations should be proposed. In the case of denying future recommendations, the time interval which is set to regularly check on user interactions will be disabled thus preventing the system from evaluating user behaviour based on the heuristics proposed in Section 5.2.3. This also prevents the modal dialogue from appearing entirely.

Should the user decide to reenable the automatic recommendations, a similar entry to the one in Figure 6.3, which is present on the modal dialogue of automatically recommended changes, can be found in the settings menu of the A+CHIS which is represented by a cogwheel on the bottom left corner of the application.

The A+CHIS' methodology of automatically adapting the LOD of visualisations is intended to support the user in maximizing their efficiency with the system. The automatic adaptations aim to provide beneficial input to further enhance the LODs the user might already have chosen manually while also respecting the user's preferences and choices. The recommendations made by the system are presented in a noninvasive manner and provide a useful preview to help the user decide if an adaptation seems advantageous or not.

CHAPTER 7

Conclusion and Future Work

Concluding this work, the A+CHIS is a consumer health information system which introduces an adaptive and interactive perspective to inferring information from medical documents. The A+CHIS allows users and patients to gather data on specific topics of interest without having to read through the entire document by visualising the most important parts while simultaneously applying topic modelling in order to highlight chapters and topics and group essential information together [1]. These visualisations can then be interacted with to enhance the search for additional information on certain keywords [4]. All of these features enable users to infer the information they deem important and interesting. Throughout this entire searching and interaction process, multiple ways of adapting the visualisations to fit the user's level of visual literacy can be made. This is done to maximize the information gain per visualisation as some visualisation might be too complex or incomprehensible for some users. Giving these visualisations multiple LODs allows them to take on a form which can fit most types of users regardless of their level of visual literacy. A change in the LOD of a specific visualisation can be induced manually by working with a matrix which entails all available visualisations and their respective LODs or automatically by letting the system decide and evaluate which LOD might be best fitting given the current behaviour and interaction history of the user with the system.

Many improvements or adjustments could still be made to this work. The most straightforward change is the extension of the currently available LODs to feature a fourth or even fifth tier which allows any type of visualisation to be even more complex or simple in order to cover a broader spectrum of advanced or inexperienced system usage. Additionally, more view components that are part of the A+CHIS could receive a rework to provide the possibility of being changed to a different LOD and then considered in the manual as well as automatic adaptation process. Furthermore, the rule-based approach for the automatic adaptation recommendations could be refined or even changed to a completely different approach which would have a better success at estimating and evaluating user behaviour to find a fitting LOD for certain visualisations.

Similarly, the heuristics for the automatic adaptations related to the number of mouse clicks per visualisation could be enhanced to cater to different view components individually since different types of visualisations might not see the same frequency of usage and require various expected degrees of mouse interactions or clicks. For example, the tilebar might not see as many mouse clicks as the topic cloud or word cloud since its main purpose as a tooltip is to indicate the number of occurrences of a term inside the document by hovering over it. This could mean that the tilebar will ultimately see more mouse movement than mouse clicks due to the nature of this type of visualisation. By specifying unique rules for every view component regarding the observed number of mouse clicks, this potential problem could be circumvented.

Another major improvement to the automatic adaptations is to include changes in the LODs of the exploration view's visualisations in the observation sessions. Tracking the adaptations made both manually and automatically could be used to further refine the automatic adaptation process. This would also enable the manual adaptations to overrule automatic recommendations even more by checking whether the user swaps back to previous LODs after having received a recommendation for changing it or denies such changes entirely over multiple instances of recommending them. The heuristics themselves could even be improved and receive new rules by recognizing user behaviour made by manual adaptations.

Lastly, the A+CHIS already features a questionnaire for first-time users in order to receive a first-hand impression of their level of visual literacy. This questionnaire presents the users with different data graphs and visualisations and asks specific questions about the data shown on these graphics in a set timeframe to get a better understanding of the user's ability to properly infer the correct information shown on these graphics. The user's level of visual literacy is finally represented by a so-called "visual literacy score". This visual literacy score can be used to preemptively adjust different settings in the system such as the LODs for certain view components to make the initial experience more suitable for the user.

The initial LODs for all view components could be adapted to fit this opening visual literacy score by setting hard-coded thresholds, similar to the automatic adaptation process, in which each visual literacy score can be linked to a specific LOD. At the application's first usage, all view components will then have their complexity set to the determined LOD representing the visual literacy score of its user.

Bibliography

[1] L. Shao, S. Lengauer, H. Miri, M. A. Bedek, B. Kubicek, C. Kupfer, M. Zangl, B. C. Dienstbier, K. Jeitler, C. Krenn, *et al.*, "Visual document exploration with adaptive level of detail: Design, implementation and evaluation in the health information domain.," in *VISIGRAPP (3: IVAPP)*, pp. 133–141, 2023. → [p1], [p3], [p6], [p9], [p10], [p19], [p22]

[2] A. Burns, C. Xiong, S. Franconeri, A. Cairo, and N. Mahyar, "How to evaluate data visualizations across different levels of understanding," in *2020 IEEE Workshop on Evaluation and Beyond - Methodological Approaches to Visualization (BELIV)*, pp. 19–28, 2020. → [p1]

[3] D. K. Ahern, "Challenges and opportunities of ehealth research," *American Journal of Preventive Medicine*, vol. 32, no. 5, Supplement, pp. S75–S82, 2007. Critical Issues in eHealth Research. → [p2]

[4] M. Tytarenko, L. Shao, T. W. Rutar, M. A. Bedek, C. Krenn, S. Lengauer, and T. Schreck, "Hierarchical topic maps for visual exploration and comparison of documents," The Eurographics Association, 2024. → [p3], [p7], [p8], [p22]

[5] S. Lengauer, M. A. Bedek, C. Kupfer, L. Shao, D. Albert, and T. Schreck, "Recognizing user behavior from interactions for adaptive consumer information systems.," in *IMET*, pp. 23–26, 2023. → [p3], [p6], [p13], [p15]

[6] T. Schreck, D. Albert, M. A. Bedek, K. Horvath, K. Jeitler, B. Kubicek, T. Semlitsch, L. Shao, and A. Siebenhofer-Kroitzsch, *Adaptive Visualization of Health Information Based on Cognitive Psychology: Scenarios, Concepts, and Research Opportunities*, pp. 165–195. Cham: Springer International Publishing, 2023. → [p3]

[7] M. Golemati, C. Halatsis, C. Vassilakis, A. Katifori, and G. Lepouras, "A context-based adaptive visualization environment," in *Tenth International Conference on Information Visualisation (IV'06)*, pp. 62–67, 2006. → [p3]

[8] J. Hussain, A. Ul Hassan, H. S. Muhammad Bilal, R. Ali, M. Afzal, S. Hussain, J. Bang, O. Banos, and S. Lee, "Model-based adaptive user interface based on context and user experience evaluation," *Journal on Multimodal User Interfaces*, vol. 12, pp. 1–16, Mar 2018. → [p3], [p13], [p19]

[9] D. Lewis, B. L. Chang, and C. P. Friedman, *Consumer Health Informatics*, pp. 1–7. New York, NY: Springer New York, 2005. → [p4]

[10] I. Ramsey, N. Corsini, M. D. Peters, and M. Eckert, "A rapid review of consumer health information needs and preferences," *Patient Education and Counseling*, vol. 100, no. 9, pp. 1634–1642, 2017. → [p4]

[11] A. Locoro, W. P. Fisher, and L. Mari, "Visual information literacy: Definition, construct modeling and assessment," *IEEE Access*, vol. 9, pp. 71053–71071, 2021. → [p5]

[12] R. Lengler, "Identifying the competencies of 'visual literacy' - a prerequisite for knowledge visualization," in *Tenth International Conference on Information Visualisation (IV'06)*, pp. 232–236, 2006. → [p5]

[13] E. Danovaro, L. De Floriani, P. Magillo, E. Puppo, and D. Sobrero, "Level-of-detail for data analysis and exploration: A historical overview and some new perspectives," *Computers Graphics*, vol. 30, no. 3, pp. 334–344, 2006. → [p5]

[14] P. Godfrey, J. Gryz, and P. Lasek, "Interactive visualization of large data sets," *IEEE Transactions on Knowledge and Data Engineering*, vol. 28, no. 8, pp. 2142–2157, 2016. → [p5]

[15] J. Baumgart, U. Viegener, and C. Pohl, *Den Diabetes im Griff: Ein Handbuch für Patientinnen und Patienten mit Diabetes mellitus typ 2.* AOK-Bundesverband, Berlin, 2021. → [p6]

[16] W. Cui, Y. Wu, S. Liu, F. Wei, M. X. Zhou, and H. Qu, "Context preserving dynamic word cloud visualization," in *2010 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 121–128, 2010. → [p7]

[17] C. Krenn, T. Semlitsch, C. Zipp, S. Lengauer, L. Shao, T. Schreck, M. Bedek, C. Kupfer, D. Albert, B. Kubicek, A. Siebenhofer, and K. Jeitler, "Customization options in consumer health information materials on type-2 diabetes mellitus—an analysis of modifiable features in different types of media," *Frontiers in Public Health*, vol. 12, 2024. → [p11]

[18] C. Inibhunu, S. Langevin, S. Ralph, N. Kronefeld, H. Soh, G. A. Jamieson, S. Sanner, S. W. Kortschot, C. Carrasco, and M. White, "Adapting level of detail in user interfaces for cybersecurity operations," in *2016 Resilience Week (RWS)*, pp. 13–16, 2016. → [p11], [p20], [p21]

[19] L. Alpay, J. Verhoef, B. Xie, D. Te'eni, and J. Zwetsloot-Schonk, "Current challenge in consumer health informatics: Bridging the gap between access to information and information understanding," *Biomedical Informatics Insights*, vol. 2, p. BII.S2223, 2009. PMID: 20419038. → [p13]

[20] P. A. Akiki, A. K. Bandara, and Y. Yu, "Engineering adaptive model-driven user interfaces," *IEEE Transactions on Software Engineering*, vol. 42, no. 12, pp. 1118–1147, 2016. → [p15]

List of Figures