UTRECHT UNIVERSITY

DEPARTMENT OF INFORMATION AND COMPUTING SCIENCES

GAME AND MEDIA TECHNOLOGY - MASTER THESIS

MovieWall: A Novel Interface Concept for Visual Exploration of Large Movie Collections

Author
Marij Nefkens

Supervisor
Dr. Wolfgang Hürst

Second Examiner
Prof. Lynda Hardman

ICA-3927679

September 22, 2017

Universiteit Utrecht

Abstract

Streaming online video content, including movies, has become extremely popular. To assist users with choosing a movie, interfaces of state-of-the-art movie streaming services generally present the user with a limited number of options selected by their recommendation systems. This allows the user to make a choice without spending too much time or effort. However, with this approach they generally lack a convenient way to explore the entire collection, without going in a very specific direction. If a user simply wants to look around, without a specific goal in mind, current interfaces do not facilitate this very well. To complement state-of-the-art movie browsing interface, we present the interface concept we call the *MovieWall*. It facilitates the exploratory and engaging exploration of large movie collections on mobile touch-based devices. The interface essentially consists of a large grid of movie posters, controlled by standard touch gestures. In this thesis, we verified the potential, feasibility and usefulness of this concept through a pilot study and a consecutive, more elaborated user study. The result showed that users were able to handle this new concept, and were comfortable with browsing relatively small images and with a lot of content on their screen at once. The subjective feedback was primarily positive. Besides testing the usefulness of the interface concept, we also compared different arrangements of the movies in the wall: randomly distributed and (semi-)clustered. Most users prefer the movies in the wall being clustered rather than randomly distributed, but that a random arrangement did not perform significantly worse in terms of user experience or satisfaction of movies encountered. Finally, we present a demo interface, improved according to the findings in the two user studies.

Acknowledgements

I would like to thank my supervisor dr. Wolfgang Hürst for his extensive guidance, feedback and enthusiasm throughout this thesis project. I would also like to thank my second supervisor prof. Lynda Hardman for her valuable feedback and insights in the field of human-computer interaction. Furthermore, I would like to thank Bruno dos Santos Carvalhal (who wrote a thesis on a similar subject) for his input in our meetings, and his willingness to assist me whenever I asked him for help. I also want to acknowledge everyone that participated in one or both of my user studies, as this thesis would not have been possible without their help. Lastly, I want to express my gratitude to my friends and family, who provided both the necessary support and distraction during stressful times.

Contents

Introduction

This research project was aimed at verifying the feasibility, potential and usefulness of the *MovieWall* concept: a novel interface concept to facilitate engaging, visual exploration of large movie collection on mobile touch-based devices. This work contains the following:

- A scientific paper discussing the main results of this thesis.
- A demo paper presenting the last, improved version of the *MovieWall* interface, based on the findings of this thesis.
- An annotated appendix that contains any material relevant to this project, that was not discussed in the scientific paper or demo paper. This includes:
  - The literature study conducted to explore previous research done in the area of exploratory movie browsing and related areas of image and video browsing.
  - Additional material of the pilot study, including implementation details, data gathered for each individual participant and an overview of the interview outcome.
  - Additional material of the second user study, including implementation details, issues encountered, data gathered of each individual participant and all interview questions and answers.
  - Any addition implementation details of the demo interface that were not included in the demo paper.
- A section summarizing and concluding this this and giving directions for future work.

Other deliverables, not included in this thesis are the following:

- Source code and executables of the applications used in the user studies, the demo interface and associated tools.
- All data gathered in both the pilot study and the second study, i.e. log files, questionnaire data and interview notes.
- A video targeted at a broader audience, introducing the *MovieWall* interface concept and summarizing the most important results of this thesis.

Chapter 1

Scientific paper

This chapter contains the scientific paper which is the main result of this thesis. It discusses the informal pilot study that verified the feasibility of the MovieWall concept and the consecutive, detailed user study further examining its potential and the influence of different user interface settings.

MovieWall: A Novel Interface Concept for Visual Exploration of Large Movie Collections

Marij Nefkens
Utrecht University
Utrecht, The Netherlands
m.w.nefkens@students.uu.nl

September 22, 2017

Abstract

Popular movie streaming services often focus on quickly providing the user with a movie or TV show to watch and generally lack a convenient way to explore beyond the recommendations. To complement such applications, we present the *MovieWall* interface concept to explore large movie collections. A pilot study with a prototype of the interface verified its potential and feasibility, but revealed some issues: it was quite overwhelming and it lacked structure. A further, more detailed evaluation of the improved interface showed even better results. Users were comfortable browsing with relatively small images, enabling them to see a lot of content. Subjective feedback was overwhelmingly positive; 80% of the participants enjoyed using the application and even more stated they would use it again. Furthermore, we showed that most users prefer the movies in the wall being clustered rather than randomly distributed, but that a random arrangement did not perform significantly worse in terms of user experience or satisfaction of movies encountered. While this work is restricted to the domain of movie browsing, we are confident this concept can be used for the exploration of other, visual collections too.

1 Introduction

Video streaming has become increasingly popular, including movie streaming. Netflix currently has over 90 million subscribers worldwide[12] and this number is still growing. The interfaces of these streaming applications generally seem to be designed for targeted search and are mostly focused on providing a user with a movie or TV show to watch without too much time or effort spent looking around. This is typically achieved by presenting the user with a limited number of movies that are likely to fit the user's interest, selected by a recommendation system. This approach makes sense for a television interface controlled by a remote, where complex navigation is often undesirable. However, nowadays a multitude of devices are used to browse and watch video content, including smartphones and tablets. For such mobile touch-based devices, which offer new and rich interaction possibilities, the lack of a convenient way to explore beyond the recommendations seems quite limiting. Providing an alternative and engaging means for users to explore all movies on offer, would nicely complement the targeted, efficiency-based approach of modern movie browsing interfaces.

A visualization that is capable of showing a relatively large number of movies, is not uncommon in advertisements or previews for any movie related services. These often feature large tile-based layouts filled with movie posters. On a screen, a layout like this would allow us to make efficient use of the limited space available. Additionally, its use in advertisements suggests that it is a visually appealing approach, aimed at making people interested in the content. So could this concept work for an interactive, engaging interface as well?

Both the lack of functionality in state-of-the-art movie browsing interfaces and the above observation lead to the idea of what we call the *MovieWall* interface concept: a movie browsing interface for the visual exploration of large movie collections, targeted for mobile touch-based devices. The essence of the concept is rather simple; a large grid with each cell containing a movie poster (the movie wall), that can be explored using common zoom and pan touch gestures. Showing movie posters only (initially) and allowing the use of zooming and panning, results in the ability to show a large collection of movies.

This approach has some obvious disadvantages however. Showing a large number of movies can be overwhelming for the user. And the more movies on a screen, the smaller the movie posters will become and could eventually become too small for people to recognise. On the other hand, research has shown that humans are quite competent at recognising small images (otherwise referred to as thumbnails) [6, 7, 22]. However, so far this research mainly focussed on the use of isolated thumbnails or only a few at a time. Thus, these results might not translate well to situations with hundreds of small images on screen at once. Moreover, if there are many small items on a touch screen device, it can become harder to interact with as accuracy comes into play. However, allowing the user to zoom should largely solve both problems, as the user herself has control over the number of items on her screen.

To establish the potential and feasibility of the concept, an informal pilot study was conducted with a prototype implementation of the interface. After verifying this was the case, the interface was improved upon and its potential was further evaluated in a detailed user study. Additionally, the influence of different arrangements of the movies was examined in this user study.

In this work, we make the following contributions:

- Provide a novel, engaging and exploratory movie browsing interface that could complement interfaces of existing movie streaming services as it has the following advantages:
  - Ability to show a large movie collection in its whole.
  - Easily explore related movies based on movie metadata. In this case we used genres, actors, director(s) and production companies.
  - Ability to easily navigate between different parts of the collection, e.g. between genres.
- Prove the feasibility and potential of the interface concept.
- Show the influence of clustering and randomization of the movie collection on the browsing experience and behaviour of users.

The paper is structured as follows: section 2 outlines related work in the fields of movie, video and image browsing. We discuss the pilot study and the subsequent detailed user study in sections 3 and 4, respectively. Section 5 concludes the work and gives directions for future work.

2 Related Work

While movie browsing has become very relevant in commercial areas due to the rise in popularity in online streaming services, the specific topic is not one to appear much in scientific research. Low et al. [10] recently proposed an interface for exploratory movie browsing using map-based exploration. Initially, the user is presented with a few popular movies and can explore the movie collection from there. Selecting a movie results in more related movies being shown, selected based on similar data like genres and actors, which are clustered on their respective type. This approach allows users to easily find interesting links between movies. A quite different approach to exploratory movie browsing is shown in [3, 11], in which users can browse through a cloud of words (e.g. moods, emotions) which were previously extracted from the audio and subtitles of movies. The interface was perceived as interesting and fresh, but slightly overwhelming.

If we look at two related fields, image and video browsing, applications often use a two-dimensional scrollable storyboard layout (e.g. standard mobile gallery applications). Yet, this design has limitations (more elaborately discussed in [19]), mainly that it can only show a limited number of items, the more items the bigger this issue becomes. Only showing a limited number of items makes it hard to get a good overview of a collection and any potential structure. To overcome these issues, many alternatives have been proposed. We see interfaces that layout all images on a large canvas, based on similarity[13, 20, 21]. Such a visualization is often created using methods like multi-dimensional scaling (MDS), which can be computationally expensive when used globally, especially with large data sets. Moreover, project images based on the outcome of methods such as MDS, often produces overlap. Rodden et al. [16] showed that overlap made it hard for users to recognise some images as edges and other important details are regularly occluded. For touch-based devices overlap can also make it harder to select items[4]. To explore large collections, graph-based and hierarchical approaches are common, sometimes combined[2]. A hierarchical layout typically allows its user to start their search very high-level, by providing an overview of images that represent the entire collection, as seen in [8, 17]. The user can then direct their search to a region of interest, revealing more similar content. This design is quite suited for displaying large collections, the exploratory similarity-based image browsing application proposed in [8] could store millions of images. Graph-based visualizations have been proposed to allow the user to browse a collection step by step, where selecting one item typically links to similar content, such as previously seen in [10]. Besides two-dimensional visualizations, there have been some three-dimensional visualizations as well, arguing that these make better use of the available screen space and can hence show more at a glance[19]. Two prime examples of 3D visualizations are shown in the interfaces in [18] and [1, 17], which respectively map a grid of images to a horizontal ring-shape and a globe-shape. Both performed quite similar on tablet and smartphone[5, 19], and outperformed the standard two-dimensional scrollable storyboard in terms of efficiency. Notable was that they also compared both three-dimensional layouts to a zoomable grid layout. This zoomable grid, where each image was initially shown in grid-like structure and could be zoomed in on, worked well on a smartphone and outperformed the others.

In the arrangement of images or videos we often encounter some degree of sorting or clustering, which can be very useful to conveniently locate specific content[16, 10]. It allows users to quickly assess the available data and judge the relevance of entire areas instead of single items. In the area of image browsing, images are often sorted on visual features (e.g. colour). However, one downside to visual similarity sorting is shown in [15], where target images were sometimes overlooked when placed alongside similar images. In a random arrangement these images stood out more, and thus were more easily found. Visual-based sorting makes less sense for movie posters [9] as visual similarity does not guarantee conceptual similarity. To capture conceptual similarity, caption-based sorting is an option. In [15] this was even slightly preferred over visual sorting. However, caption-based arrangements rely on good image annotations. Luckily, movies generally have a lot of available metadata.

3 Pilot Study

As the foundation of the interface was built mainly on assumptions, it was important to verify the potential and feasibility of the interface concept in an early stage. Hence, a small pilot study was conducted. The goals for the pilot study are outlined in Section 3.1. The implementation of the prototype interface is described in Section 3.2. Section 3.3 explains the procedure of the experiment. Finally, the results are shown in Section 3.4 and discussed in 3.5.

3.1 Objectives

The main question to answer in this pilot study was the following: *Does the MovieWall interface concept have potential as an alternative for exploratory movie browsing?*

The first requirement is that users need to be able to handle the interface. So, *are users able to (easily) interact with the interface?* Also, *are users not too overwhelmed by it?* It's important that users are not negatively influenced by either.

The second important aspect of the concept is that it should provide an engaging experience. In order to do so, users need to enjoy interacting with the interface to some degree, which is another important variable to assess.

Furthermore, we wanted to analyse the interaction behaviour of participants. *How do they use the available functionalities? What zoom levels do they use?* The latter could give more insight in what size images people are comfortable browsing. If this turns out very large, perhaps the idea of presenting a lot of movies does not have as much potential, as users are not comfortable with so many items on their screen.

Figure 1: Left: screenshot of the initial and at the same time most zoomed out state of the prototype application. Right: most zoomed in state of the prototype application. In this case a movie is selected and thus a panel is shown containing the details of the selected movie. All movie posters were downloaded from www.themoviedb.org and used under the fair usage policy.

Finally, if there is potential, *what issues need to be addressed in order to improve the interface?*

3.2 Implementation

The prototype application is developed for Android and designed for tablet use, with a reference resolution of 1920 × 1200 pixels (16:10). All images and data used were downloaded from the The Movie Database[1] (TMDb), either during runtime or pre-stored, using TMDb's open API.

For the pilot study, a first and basic version of the interface concept was implemented. It features the core element of the concept: a large wall of 8640 movie posters (144×60 movies), which are the 8640 most popular movies in The Movie Database. The user can change both her position in the grid and number of movie posters on her screen by zooming and scrolling.

The initial, and at the same time most zoomed out state of the movie wall, is shown in Figure 1 (left). In case of the reference resolution, all movies are shown on the screen, each poster image with an average size of 13.3 × 20 pixels.

Based on earlier research on recognition of small images[22], we assumed it would be near impossible for users to recognise anything at this point and that this would encourage them to use zoom in and stop at a zoom level on which they could comfortably browse. The most zoomed in version of the interface is shown in Figure 1 (right), in this case the movie posters are 248 × 372 pixels on screen and should be easily recognisable.

Selecting a movie marks its poster in the grid to indicate the current focus, and opens a panel with the main movie details like actors and plot as shown in Figure 1 (right). Instead of taking up the entire screen space, the panel only occupies part of the available screen space. This allows users to still interact with the movie grid while the panel is open, but with a more limited space to do so.

The arrangement of movies in the grid is random. We did not expect the sorting of the movies to significantly influence the feasibility or potential of the concept. Furthermore, we hypothesized that a random arrangement could actually provide an interesting environment which allows for serendipitous discoveries.

3.3 Procedure

The experiments took place in a controlled environment with an observing party. Participation was entirely voluntary and not reimbursed. Prior to the experiment, participants filled out a consent form and received a short introduction to the experiment and interface. The test device used was a Samsung Galaxy Tab A with a screen resolution of

[1] TMDb website: https://www.themoviedb.org/

1920 × 1200 pixels.

Prior to the actual trial, the participant was asked to play around with the application for a short while (max. 2 minutes). This was both to familiarize the participant with the interface and to observe how the participant interacted with the interface when using it for the first time. Afterwards, the participant was asked to read a list of all functionalities of the interface and state if they missed anything.

Next, the participant was asked to use the application with the following task (max. 5 minutes): *"Assume you are about to board a long flight, so you will be offline for a couple of hours. Hence, you want to download a couple of movies to watch when you are in the air. Now use the interface to search for, let's say 3-5 movies that you would probably download."*

Afterwards, the participant filled out a questionnaire and a semi-structured interview was conducted about the answers given in the questionnaire and the observed behaviour while conducting the task.

During both trials, the participant was observed and each interaction was logged by the application.

3.4 Results

Ten people participated in the pilot study, nine male and one female, ages 21-26 (average of 23.2 years old). Each of them watches movies on their computer, seven of them also on television, no one on tablets.

Are users able to handle the interface?

Generally all functionalities of the interface were quite clear to the participants. For two it took a while to realize they could zoom in initially.

Participants were asked to rate a few statements on a five-point Likert scale about their experience with the interface. These statements and the ratings given are shown in Table 1. Every participant thought the interface was easy to use (S1), and most also thought it was intuitive to interact with (S2). Participants stated they thought the interface worked fluently and according to expectation. One person was neutral about this statement, as this participant was unsure what to do with the application initially.

Most participants stated they were (slightly) overwhelmed by the interface (S3), mostly at the beginning. Only one participant disagreed with this statement. The main reasons for this were the vast number of movies, the lack of a starting point and the lack of structure. Three out of the nine participants that stated they were overwhelmed, viewed this as a clearly negative experience.

Table 1: Frequency of each rating given for each statement from the questionnaire.

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
S1	0	0	0	2	8
S2	0	0	1	3	6
S3	0	1	2	6	1
S4	0	2	3	4	1

S1: I thought the interface was easy to use.
S2: I thought interacting with the interface was intuitive.
S3: I was overwhelmed by the interface.
S4: I thought browsing the interface was enjoyable.

Do users enjoy using the interface?

Participants were divided about enjoying the browsing experience (S4). Participants who liked it, mentioned that they appreciated casually browsing around and that they discovered new and unexpected movies. Participants who did not (really) like it, stated that it was too overwhelming, that they encountered too many uninteresting movies or that they did not like the lack of options to direct their search.

Most of the participants stated they would find such an interface a useful complement to existing movie streaming applications, for example Netflix, especially with the addition of more structure or another means to direct their search to some degree. Participants could see themselves use it when not knowing or caring what to watch, or when picking a movie with multiple people, for example friends or family.

How do users interact with the interface?

Observations during the experiment and the log data indicate that the participants used the zoom function mainly to get to a comfortable level and generally remained on this level. This suspicion was confirmed in the interviews. The distribution of zoom levels used per participant are shown in Figure 2.

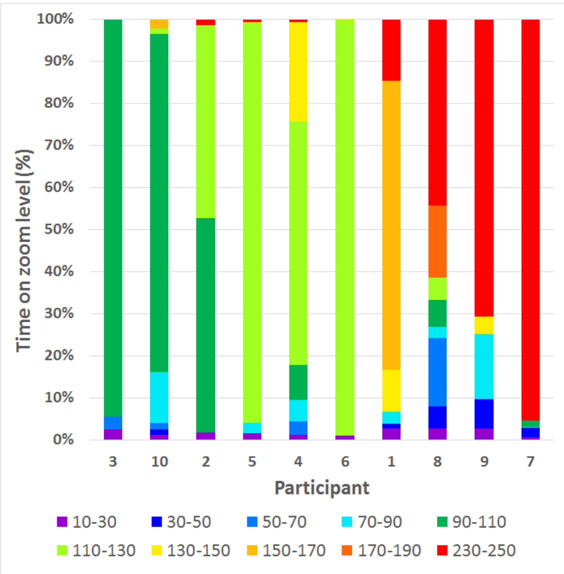

Figure 2: Distribution of zoom levels used per participant. The zoom levels ranges indicate the width of the images shown in pixels.

3.5 Discussion

All in all, the pilot study results indicate that the concept has potential, but there are some clear issues that need to be addressed in order to improve the interface. Initially, the interface was quite overwhelming for most users, which seemed to be primarily caused by the fact that the wall was completely zoomed out at the start. Starting with a larger initial zoom level could diminish this feeling. Another cause for the overwhelmed feeling was the lack of structure. While the random arrangement of movies helped users to be more open-minded and was often appreciated in the beginning, participant generally wanted more structure (after a while). There was no means for users to direct their search to some degree, which was frustrating for some. Sorting the movies based on metadata is a possible improvement here. Yet, we should keep in mind that the random arrangement did allow for more serendipitous discoveries, and arranging the movies in clusters might negate this effect.

4 Detailed User Study

The pilot study was rather small and informal, but allowed us to identify issues in the initial design and areas that needed more investigation, such as the movie arrangement. To further evaluate the usefulness of the concept, the interface was improved and a second, more detailed user study was conducted. The concrete goals of this user study can be found in the following section, 4.1. The improvements in the implementation of the interface are listed in Section 4.2. The procedure of the experiments is described in Section 4.3. Finally, the results of the experiments and its discussion can be found in Sections 4.4 and 4.5.

4.1 Objectives

The pilot study identified two potential issues of the prototype interface, namely:

1. It was overwhelming, likely caused by (a) the number of movies, (b) the initial zoom level and (c) the lack of structure.
2. There was a lack of structure, giving users no control over what content they came across, causing some users to encounter a lot of uninteresting movies.

We hypothesize that these issues can be resolved adapting the following:

a. Reducing the number of movies in total.

b. Initially showing less movies on screen.

c. (Semi-)Sorting the collection and/or adding filter functionality.

While the pilot study suggests a preference for a sorted arrangement over a random one, it needs further investigation. Arranging movies in clusters based on genre provides a means for users to direct their search and might make the interface less overwhelming. However, this is likely at the cost of serendipity. A potential middle ground is to cluster the movies, but introduce a certain degree of randomness. This approach can provide structure for the user, but prevent a narrow vision. To test these assumptions, we compared three different arrangements of movies in this study:

1. Random (R): Same as in the pilot study; all movies are randomly distributed.
2. Clustered (C): All movies are strictly clustered on genre, with the most popular movies in the center of the cluster. A possible arrangement is visualized in the left image in Figure 3.
3. Semi-clustered (CR): Same as (C), but with a certain degree of randomization in the outer areas of each cluster and within the popularity sorting. A possible arrangement is visualized in the right image in Figure 3.

Another option to provide the tools to direct the exploration process, is to allow users to filter on metadata like genre or actors. In contrast to clustering, this is an optional tool for structure. Moreover, it can be combined with any arrangement.

All in all, the following was investigated in this user study:

- The general usefulness and acceptance of the interface concept.
- The usefulness of filter functionality.
- The influence of different arrangements of the movie collection, more specifically:
  - The *interaction with the movie collection*; i.a. the number of movies selected while browsing the interface.
  - The *satisfaction of movies encountered* while browsing the interface. We suspect this is higher for the clustered arrangements, as the users are in control of what content they encounter. However, serendipitous discoveries could increase the satisfaction for the random arrangement as well.
  - The *intrinsic motivation* of users to use the interface with each arrangement. We assume that users will like the (semi-clustered) arrangements better than the random one, and thus have a higher intrinsic motivation to use these.

4.2 Implementation

The application used in the second study was an improved version of the prototype tested in the pilot study. The most important changes are listed in this section.

The grid contains less movies, 2160 in total (72×30). At the smallest zoom level each image is approximately 26×39 pixels. This decrease in movies potentially makes the interface less overwhelming. Furthermore, by decreasing the total number of movies, the average popularity of the movies increases and we expect this presents a more interesting and relevant collection of movies to the user. Additionally, the initial number of movies on the

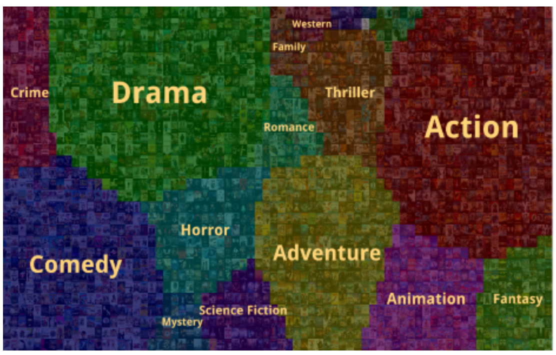

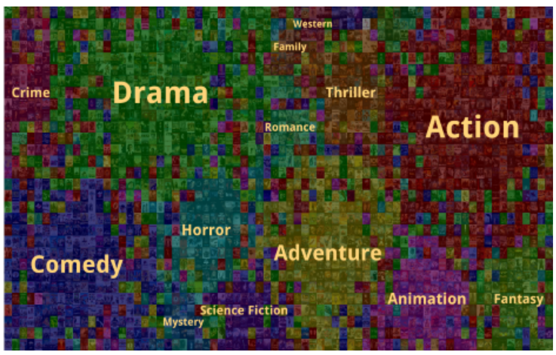

Figure 3: Visualization of a clustered (left) and a semi-clustered arrangement (right).

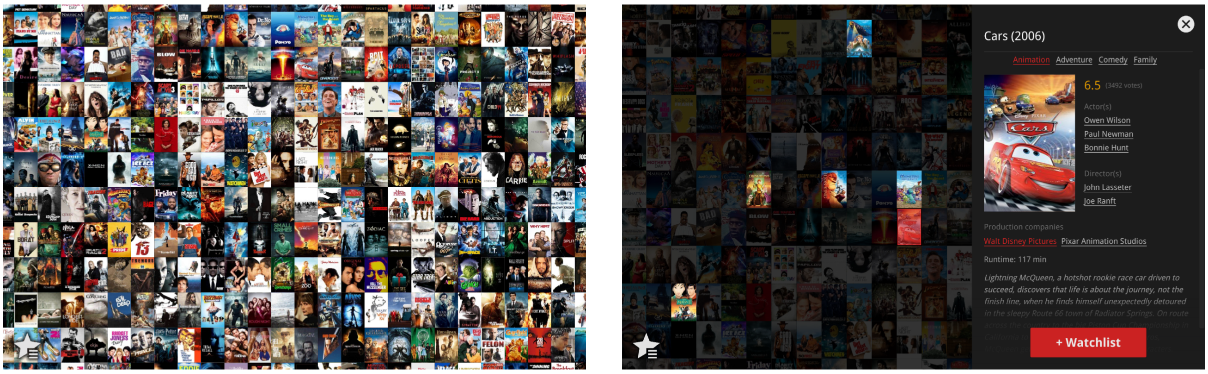

Figure 4: Left: the initial state of the improved interface. Right: The movie wall with multiple filters applied: *Animation* (genre) and *Walt Disney Pictures* (company), which highlights movies with the genre *Animation*, produced by *Walt Disney Pictures*. All movie posters were downloaded from www.themoviedb.org and used under the fair usage policy.

screen is greatly decreased to make the interface less overwhelming. The new initial state of the application is shown in Figure 4 (left).

Filters were added so users can look for movies with similar metadata as the currently selected movie. Movie details in the panel can be selected to highlight all movies with, for example, a common actor. Filters can be combined so the user can easily search for movies based on multiple criteria, without having to enter complex queries. An example is shown in Figure 4 (right). The following metadata can be applied as a filter:

- Genres (up to five)
- Top actors (up to three)
- Top directors (up to two)
- Production companies (up to six)

The clustered and semi-clustered arrangements were added. Moreover, when zoomed out far enough, labels are shown indicating the location of each genre cluster, helping the user identify the structure of the collection.

4.3 Procedure

The experiments took place in a controlled environment with an observing party. Participation was entirely voluntary and not reimbursed. Prior to the experiment, participants filled out a consent form, a list of general background questions and received a short introduction to the experiment and interface. The same device (Samsung Galaxy Tab A) was used as during the pilot study.

Each participant tested all three different arrangements. The order in which the arrangements were presented to each participant were rotated between participants, resulting in a total of six different orders. To familiarize the participants with the application prior to the actual experiment, they were asked to use it for a short while (max. 2 minutes) with the task to figure out all functionalities. Afterwards, the participants were informed about anything they possibly missed.

For the upcoming three trials, the participants received the following scenario: *"Imagine you're going on trip next week, you're going to visit South-Africa and subsequently Brazil. After your trip you're going back home, meaning you'll have three long flights ahead of you. During these flights you want to watch some movies. There is no WiFi on the plane, so you'll have to download the movies in advance. Assume that you're at the airport gate for your first/second/third flight and you have some spare time. Find some movies to download that you might like to watch during your the flight and add them to your watchlist."*

For each different arrangement the participant used the application for three minutes, starting with the same arrangement as used while familiarizing with the application. Prior to each new trial, the new arrangement of movies was globally explained to the participant. In case the participant started with the random arrangement, the cluster feature was explained prior to the second trial. During each trial, all interaction was logged by the application. After each trial, the participant rated several statements on a seven-point Likert scale, consisting of the interest/enjoyment subscale of the Intrinsic Motivation Inventory (IMI)[14] and specific questions about movies encountered (listed below Figure 6). A seven-point Likert scale was a requirement for the IMI and was used for the other statements for the sake of consistency.

After completing the three trials, a semi-structured interview was conducted with the participant concerning their opinion on the different arrangements and the interface in general.

4.4 Results

Thirty people participated in the study, of which twenty-two were male and eight female, ages 20 to 29 (on average 23.3 years old). Ten of them watched movies on mobile devices, six on tablets and five on smartphones. Twenty-two participants used movie streaming services (e.g. Netflix), of which nineteen on their computer, nine on television, nine on smartphones, four on tablets and one using the Xbox. Five participants had also participated in the pilot study.

General usefulness and acceptance

As with the pilot study, all functionalities were quite clear to the participants. As positive aspects of the interface, they mentioned in the interview afterwards it worked smoothly (8 subjects) and was easy/intuitive to use (5). In contrast to the pilot study, not all participants (5) realized it was possible to zoom. This was probably caused by the difference in initial zoom level, where during the second study it was deliberately set to a level where we assumed users could recognise what was pictured in most of the posters without being forced to zoom in first. Yet, none of these five commented on the small image size (initially 77×115 pixels per poster), suggesting they were not bothered by it.

Looking at the zoom levels, the suspicion that users are comfortable browsing with relatively small images seems further confirmed. Figure 5 shows the percentage of time on each zoom level of all participants, of both the pilot study and the random arrangement of the second study. In the second study we see that over all arrangements, roughly two-thirds of the time a zoom level between 50-150 was used. This ranges from an image width of 50×75 pixels to 150×225 pixels. With the test device used, this roughly translates to between 60 and 610 movies on screen at once with the panel closed. While there are clear similarities in zoom levels between the two studies, participants of the pilot study used larger zoom levels on average. We hypothesize that this is mainly caused by difference in initial zoom level. In the pilot study, the initial image size was on average 13.3×20 pixels, which is so small that users were forced to zoom in to be able to recognise the movie posters better. We hypothesize that this made users zoom in more than needed to recognise the posters or even browse comfortably, possibly caused by the fact that people are simply used to larger image sizes. In contrast, in the second study the initial image size was roughly 77×115 pixels, a zoom level on which we expected users to be able to recognise most content. The data suggests this was a correct assumption, as on average around 20% of the time was spent on this zoom level.

In the interview following the experiment, participants were mostly positive about the interface. When asked for their opinion on the interface, comments included: *"I saw a lot of movies (in a short time)"* (10); *"it works smoothly"* (8); *"I liked the filters"* (8); *"I encountered a lot of unexpected/interesting movies"* (7); *"I liked the use of movie posters"* (6). The main negative comment was that one could only judge a movie by its poster (at first sight) (5). Two participants mentioned they felt like this caused them to only pick familiar movies. When asked whether they enjoyed using the interface, 24 participants gave an affirmative answer, 4 were neutral about the statement and 2 said no. Both participants who did not enjoy using the interface stated that they did not like looking for movies in general. When asked whether they thought they would use an interface like this again, 26 participants said they probably would. Many stated that they would use it if they had no idea

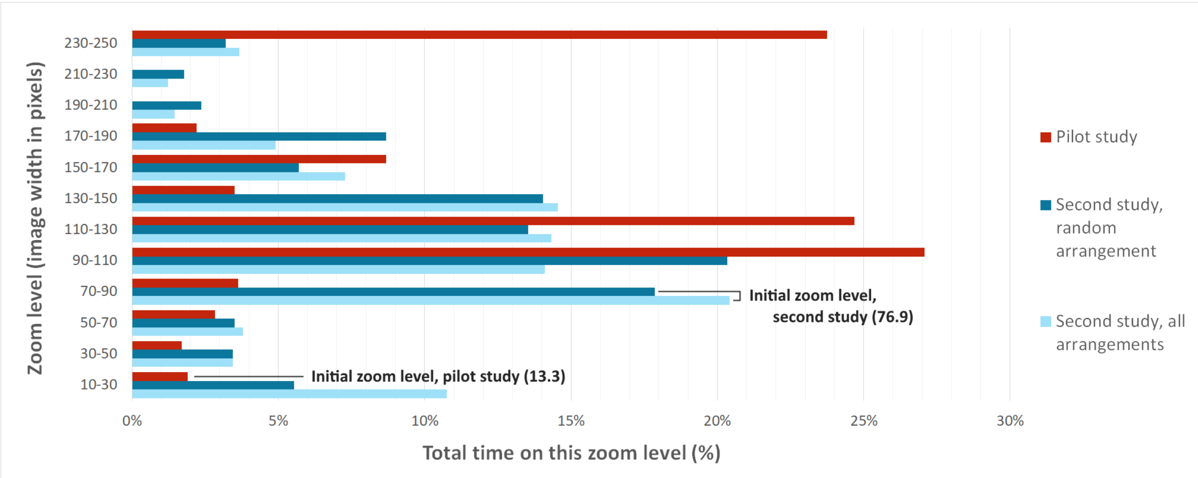

Figure 5: Comparison of zoom levels used between the pilot study, the random arrangement of the second study and the second study in total.

what to watch (10). Others stated they would use it to find a bunch of movies, for example before a vacation or flight (3), or when they wanted to watch a movie they don't know (3). Two participants even stated they would use it instead of existing movie streaming interfaces, if it contained all the movies that streaming service has to offer.

When asked what participants would change or improve upon, the most common answer was the addition of a search function (14). Many participants mentioned they appreciated having a "starting point" in the wall, for example a movie they liked and from where they could search for more related movies, which a search function could help accommodate. Another requested feature was to keep the filters active while switching or deselecting movies (9).

Filters

The filter functionality was frequently used. One or more filters were used in 66 of the 90 trials (73.3%) and only two participants did not use them in any of their trials. The number of times each kind of filter was used, is shown in Table 2. Genre filters were used most, which we assume is the most common search criteria for finding a movie to watch. Yet, lesser common search criteria like director or production company were also used relatively frequently. This indicates a desire to find related movies based on diverse sorts of criteria. Participants also mentioned they would like to be able to search on more kinds of metadata like collections (e.g. Star Wars), script writers or story origin.

Table 2: Number of times each kind of filter was used.

	Genre	Cast	Director	Company
R	79	21	0	15
C, CR	47	48	16	22
Total	126	69	16	37

Comparison of different movie arrangements

To compare the interaction of users with the movie collection between the different arrangements, the number of (unique) movies inspected during each trial and the number of movies in the watchlist at the end of each trial were recorded. We expected that both these variables could be strongly influenced for the second and third trials, since in each trial the same data set was used. Thus, only the results of the first arrangement used by each participant were considered for this comparison. Table 3 contains the difference in inspected movies and Table 4 the difference in movies added to the watchlist between the three arrangements. When performing a One-Way ANOVA, we see a significant difference in the number of movies inspected, $F(2,27) = 4.429$, $p = 0.022$. There are significantly more movies inspected with the random arrangement than with the semi-clustered arrangement. No significant results were found when inspecting R vs. C or CR vs. C ($p = 0.223$, and $p = 0.098$ respectively). The differences in movies added to the watchlist between the different arrangements are not significantly different; $F(2,27) = 0.788$, $p = 0.465$.

Table 3: Descriptives for the number of movies inspected during the first trial of each participant.

	N	Mean	SD	Min	Max
R	10	10.9	3.51	4	15
C	10	9.3	2.83	5	14
CR	10	7.1	2.079	3	10
Total	30	9.1	3.188	3	15

Table 4: Descriptives for the number of movies in the watchlist at the end of the first trial of each participant.

	N	Mean	SD	Min	Max
R	10	6.3	3.917	1	11
C	10	5.9	2.331	3	10
CR	10	4.8	1.476	3	8
Total	30	5.67	2.746	1	11

After each trial participants were asked to rate seven statements on a 7-point Likert scale. The ratings given by the participants are shown in Figure 6. None of these were significantly different between the different arrangements. Besides specific questions about movies encountered and picked during the trials, the participants were asked to rate the seven statements of the interest/enjoyment subscale of the Intrinsic Motivation Inventory after each trial. These results are shown in Figure 7. The intrinsic motivation of users was not significantly different between the different arrangements.

In contrast to the previous results, the interview showed a clear favourite arrangement. Figure 8 shows the preferred and least preferred arrangement of each participant. It shows that one-third of the participants preferred either the clustered or semi-clustered version; of those most mentioned they did not really see the difference between the two arrangements (7 subjects). Only one person specifically preferred the semi-clustered version and ten the strictly clustered version. The random version was least liked, with two-thirds (20) of the

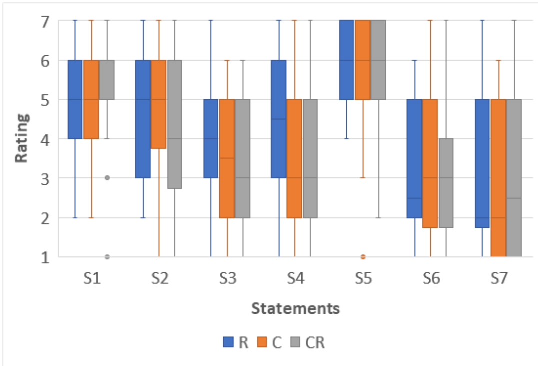

S1: I saw a lot of interesting movies
S2: I saw a lot of movies that didn't interest me
S3: I encountered a lot of movies that I didn't expect to see
S4: I thought it was hard to find movies I wanted to watch
S5: I was satisfied with the movies I picked for my fictional flight
S6: The movies I picked for my fictional flight were different from what I expected to pick
S7: The movies I picked for my fictional flight were different from what I normally watch

Figure 6: Visualization of the ratings given per arrangement by the participant for each statement in the questionnaire. The ratings were given on a 7-point Likert scale, from "not at all true" (1) to "very true" (7), with "somewhat true." (4) in the middle.

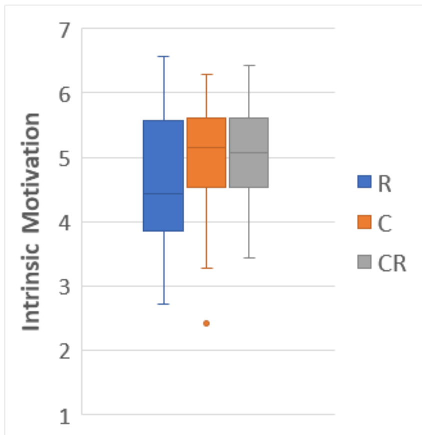

Figure 7: Visualization of the intrinsic motivation values per arrangement.

participants stating it was their least favourite. However, five participants liked the random version best. Three participants had no preference, all of them stated they would use the different arrangements in different situations (*"without specific intent I'd use the random version, otherwise one of the clustered ones"*).

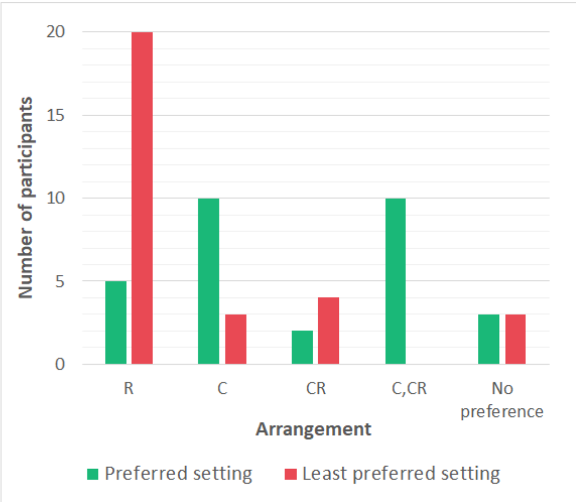

Figure 8: Number of participants that preferred and least preferred each arrangement.

The interviews with the participants indicated that the randomization in the semi-clustered version mostly stayed unnoticed. Only five participants explicitly mentioned they were surprised with the placement of certain movies, of which three preferred the strictly clustered version. While mostly unnoticed, participants did inspect and pick randomized movies. Figure 9 shows the percentage of randomized movies in the watchlist at the end of the trial for each participant. Fourteen participants had at least one randomized movie in their watchlist. Interestingly enough, two out of the three participants that preferred the strictly clustered version over the semi-clustered version because of the randomization, had randomized movies in their watchlist (participant 6 and 25).

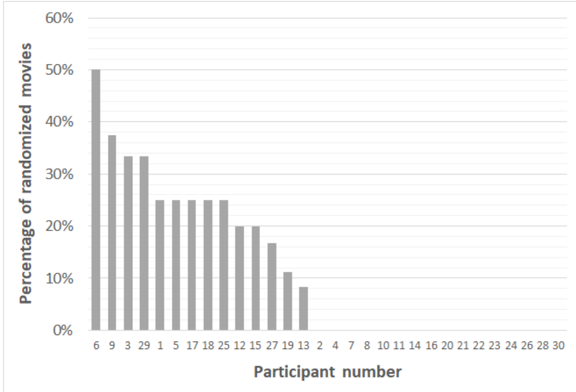

Figure 9: Percentage of randomized movies that each participant had in their watchlist at the end of the trial with the semi-clustered arrangement (CR).

4.5 Discussion

In the pilot study, we identified two potential issues of the interface: it was overwhelming and there was a lack of structure (or any other means to direct the search process to a certain extent). The results of the second study suggest that these issues have been resolved. Consequently, that seems to have lead to a more enjoyable, engaging browsing experience. When asked whether the participants enjoyed using the application, 80% of the them gave an affirmative answer, as opposed to only half of the participants in the pilot study.

While the study further confirmed the usefulness of the interface concept, the influence of the different arrangements on the tested variables was less apparent than expected. We hypothesized that the semi-clustered version could provide a nice balance between structure and allowing the user to discover unexpected content. However, there is little evidence this was the case. The issue with the semi-clustered arrangement is most likely that it opposes the expectations of users, which can be undesirable. When participants noticed the randomization in the clustered version, it was generally perceived as strange or faulty, both rather negative characteristics.

5 Conclusion & Future Work

In this work, we presented an interface concept for engaging, exploratory movie browsing we call the *MovieWall* interface concept.

One concern of this concept was the number of movies on screen could be too large and the sizes of the movie posters too small. However, participants were quite able to handle the interface. The interaction of users with the interface suggests that they were comfortable browsing with relatively small images and thus many movies on their screen at once. Previous research already showed that humans are quite proficient at recognising small images, but the fact that they are also willing to browse with such small images is new. Furthermore, the large majority stated they enjoyed using the interface and would probably use it again.

The filters added in the second version of the interface were used frequently and explicitly mentioned by users as a very useful feature. While users mostly filtered on genre, we saw a moderate use of the cast, company and director filters as well. This indicates a desire of at least some users to navigate by diverse kinds of metadata, which is often not encouraged or focussed on in state-of-the-art movie browsing applications. Some participants stated they would like even more sorts of metadata to navigate by. For future work, it would be interesting to further analyse what types of metadata people would use while browsing movies, and how to best facilitate this.

As for the arrangement of movies, the results were inconclusive. In contrary to our prediction however, the randomized arrangement did not perform significantly different than the other arrangements in terms of movies found or satisfaction about movies encountered. However, when asked what arrangement participants preferred, the strictly clustered arrangement was the clear favourite, and the random arrangement was liked least.

For future work it would be interesting to integrate this design in an existing state-of-the-art movie browsing interface, to analyse its usage. This would provide a more realistic test setting. Furthermore, it would be interesting to test the potential of the interface for different platforms such as smartphones or computers, as this work was solely focussed on tablets. Finally, we are positive this interface can be applied outside the domain of movie browsing as well, as it can essentially provide a visually appealing starting point to the exploration of any collection.

References

[1] David Ahlström et al. "A user study on image browsing on touchscreens". In: *Proceedings of the 20th ACM international conference on Multimedia*. ACM. 2012, pp. 925–928.

[2] Kai Uwe Barthel, Nico Hezel, and Radek Mackowiak. "Graph-based browsing for large video collections". In: *International Conference on Multimedia Modeling*. Springer. 2015, pp. 237–242.

[3] Nuno Gil et al. "Going through the clouds: search overviews and browsing of movies". In: *Proceeding of the 16th International Academic MindTrek Conference*. ACM. 2012, pp. 158–165.

[4] Ai Gomi and Takayuki Itoh. "Mini: A 3d mobile image browser with multi-dimensional datasets". In: *Proceedings of the 27th annual ACM Symposium on Applied Computing*. ACM. 2012, pp. 989–996.

[5] Marco A Hudelist, Klaus Schoeffmann, and David Ahlstrom. "Evaluation of image browsing interfaces for smartphones and tablets". In: *Multimedia (ISM), 2013 IEEE International Symposium on*. IEEE. 2013, pp. 1–8.

[6] Wolfgang Hürst et al. "Keep moving!: revisiting thumbnails for mobile video retrieval". In: *Proceedings of the 18th ACM international conference on Multimedia*. ACM. 2010, pp. 963–966.

[7] Wolfgang Hürst et al. "Size matters! how thumbnail number, size, and motion influence mobile video retrieval". In: *Advances in Multimedia Modeling* (2011), pp. 230–240.

[8] Yanir Kleiman et al. "Dynamicmaps: similarity-based browsing through a massive set of images". In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM. 2015, pp. 995–1004.

[9] Yanir Kleiman et al. "Toward semantic image similarity from crowdsourced clustering". In: *The Visual Computer* 32.6-8 (2016), pp. 1045–1055.

[10] Thomas Low et al. “Exploring Large Movie Collections: Comparing Visual Berrypicking and Traditional Browsing”. In: *International Conference on Multimedia Modeling*. Springer. 2017, pp. 198–208.

[11] Pedro Martins, Thibault Langlois, and Teresa Chambel. “MovieClouds: content-based overviews and exploratory browsing of movies”. In: *Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments*. ACM. 2011, pp. 133–140.

[12] Netflix. *Netflix website targeted for investors*. URL: `https://ir.netflix.com/index.cfm`.

[13] Zoran Pečenovió et al. “Integrated browsing and searching of large image collections”. In: *International Conference on Advances in Visual Information Systems*. Springer. 2000, pp. 279–289.

[14] Robert W Plant and Richard M Ryan. “Intrinsic motivation and the effects of self-consciousness, self-awareness, and ego-involvement: An investigation of internally controlling styles”. In: *Journal of Personality* 53.3 (1985), pp. 435–449.

[15] Kerry Rodden et al. “Does organisation by similarity assist image browsing?” In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM. 2001, pp. 190–197.

[16] Kerry Rodden et al. “Evaluating a visualisation of image similarity as a tool for image browsing”. In: *Information Visualization, 1999.(Info Vis' 99) Proceedings. 1999 IEEE Symposium on*. IEEE. 1999, pp. 36–43.

[17] Gerald Schaefer. “A next generation browsing environment for large image repositories”. In: *Multimedia Tools and Applications* 47.1 (2010), pp. 105–120.

[18] Klaus Schoeffmann and David Ahlström. “Using a 3d cylindrical interface for image browsing to improve visual search performance”. In: *Image Analysis for Multimedia Interactive Services (WIAMIS), 2012 13th International Workshop on*. IEEE. 2012, pp. 1–4.

[19] Klaus Schoeffmann, David Ahlström, and Marco A Hudelist. “3-D interfaces to improve the performance of visual known-item search”. In: *IEEE Transactions on Multimedia* 16.7 (2014), pp. 1942–1951.

[20] Grant Strong, Orland Hoeber, and Minglun Gong. “Visual image browsing and exploration (Vibe): User evaluations of image search tasks”. In: *International Conference on Active Media Technology*. Springer. 2010, pp. 424–435.

[21] Damien Tardieu et al. “Browsing a dance video collection: dance analysis and interface design”. In: *Journal on Multimodal User Interfaces* 4.1 (2010), pp. 37–46.

[22] Antonio Torralba, Rob Fergus, and William T Freeman. “80 million tiny images: A large data set for nonparametric object and scene recognition”. In: *IEEE transactions on pattern analysis and machine intelligence* 30.11 (2008), pp. 1958–1970.

Chapter 2

Demo paper

This chapter contains the demo paper, meant to be submitted to the Multimedia Modeling Conference (MMM 2018). It introduces the *MovieWall* interface concept and showcases the improved version of the *MovieWall* interface as shown in the previous chapter, based on the finding of both the pilot study and the second, more detailed user study.

MovieWall: A Novel Interface Concept for Visual Exploration of Large Movie Collections

Marij Nefkens

Utrecht University
m.w.nefkens@students.uu.nl

Abstract. With the continuous growth in video streaming and video content, it becomes increasingly important to make the available content accessible to users. In the domain of movie browsing, popular streaming applications often present the user with a limited number of recommended options. This allows users to make a choice relatively fast and effortless, but generally provide no convenient way to explore beyond the recommendations without going in a very specific direction. This approach can be limiting, especially on touch-based devices which offer rich interaction possibilities. In this work we demonstrate an interface concept we call the *MovieWall* in order to complement existing systems. It facilitates exploratory browsing through large movie collections on mobile touch-based devices. The interface essentially consists of a large grid of movie posters, controlled by standard touch gestures. Simplistic, yet visually appealing and able to display a large number of movies. User studies have verified the feasibility and usefulness of the interface.

1 Introduction

Video streaming has become extremely popular in the past years, as a result of the increasing accessibility of (fast) internet and the ever growing amount of video content, both professional and user generated. These developments make it increasingly important to find suitable ways to present the available content to the user. Interfaces of state-of-the-art movie streaming services generally use the approach of recommendations. The system decides what (small) part of the available content is presented to the user, so the user does not have to spend a lot of effort and time browsing around. This approach works well for specific situations and certain platforms, but can be limiting as there is generally no convenient way to explore beyond the recommendations without going in a very specific direction. On a television controlled by a remote, an interface requiring a lot of complex interaction is undesirable. However, nowadays a multitude of devices are used to browse and watch video content. Especially touch-based mobile devices, which have become very dominant on the market, offer rich interaction possibilities. Yet, state-of-the-art movie browsing interfaces do not differ all that much between platforms.

To complement existing movie browsing interfaces, we present what we call the *MovieWall*: an interface concept for engaging, visual exploration of large movie collections on mobile touchscreen devices. The interface essentially consists of a large grid of movies, presented by their respective movie poster, which can be navigated using standard zooming and panning touch gestures. It is a simple, yet visually appealing design, which can show a relatively large number of movies as only their posters are shown. This approach is inspired by the fact that humans are quite capable of recognising small images. In [2] it was shown that with images as small as 32×32 pixels, people were able to recognise over 80% of their content accurately. Comparable results were found in [1], when evaluating the use of small thumbnails on mobile phones, contradicting the common belief that these devices are not suitable for showing multiple small images at the same time. Furthermore, the user studies we have conducted with the *MovieWall* interface strongly indicate that users are not only able to browse relatively small movie posters, but are also comfortable doing so.

2 Interface Overview

The demo interface is implemented for Android and designed for touchscreen tablets. The device on which all user studies have been conducted, has a resolution of 1920×1200 pixels. All movie related data used, is obtained from The Movie Database (*www.themoviedb.org*), through their open API.

Movie wall The core element of the application is the movie wall: a large grid with movie posters. In the demo interface it consists of 2160 movies. This number of movies was not perceived as too overwhelming in the user studies conducted hence used for this demonstration. We are confident this number can be increased, as long as the movie collection remains interesting and relevant for the user. The movie wall can be navigated using standard touch gestures: multi-touch zooming and single-touch dragging. Completely zoomed out, the entire wall is shown to give users an idea of the size and structure of the collection, as shown in Fig. 1 (left). Completely zoomed in, the posters should be large enough so every detail is easily recognisable.

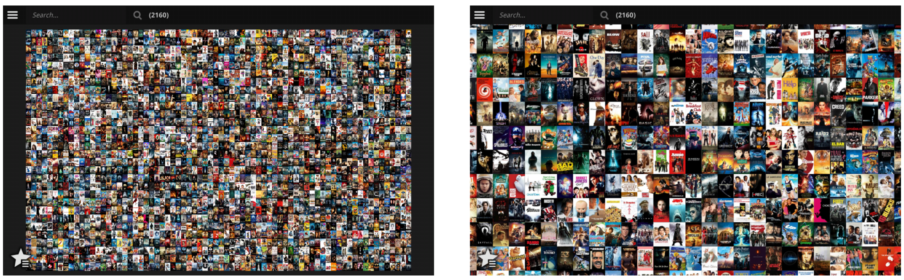

Fig. 1. Left: Completely zoomed out movie wall. Right: Zoom level at the start. In both screenshots the movies are randomly distributed over the grid.

Zoom levels The zoom level at starting point of the application is shown in Fig. 1 (right). Initially, poster images are relatively small, likely smaller than what people are used to, but still recognisable (approximately 77 × 115 pixels). User tests we conducted strongly indicate that this initial zoom level contributed to the users' tendency to browse with relatively small images. This is encouraged considering it allows users to see large amounts of content at once, which is the idea behind the interface. Roughly two-thirds of the time, users browsed with image sized between 50 × 75 pixels and 150 × 225 pixels. With the test device used in these experiments, this approximately translated to between 60 and 610 movies on screen at once.

Movie details & filters Selecting a movie in the grid opens a side panel, containing a short overview of its metadata; the movie title and its most relevant genres, cast, directors and production companies (Fig. 2, left). Each of the latter can be applied as a filter. A filter highlights every movie in the wall with the same metadata, as shown in Fig. 2. Multiple filters can be combined to allow the user to search for related movies based on various criteria; a powerful yet easy to use feature, which does not require the user to enter complex queries. The filters can be turned off by either deselecting them in the side panel or by removing them in the top bar. The user can add custom filters as well, by entering keywords in the input field in the top bar, as shown in Fig. 2 (right). In the user tests, we observed that the filters were used often by participants. Furthermore, each type of metadata provided as a filter was used to some extent.

To show additional movie details, the information button in the side panel can be clicked. These movie details are shown instead of the movie wall, as it is assumed users generally do not need to see both at the same time. Furthermore, this design creates enough space to show additional material like trailers.

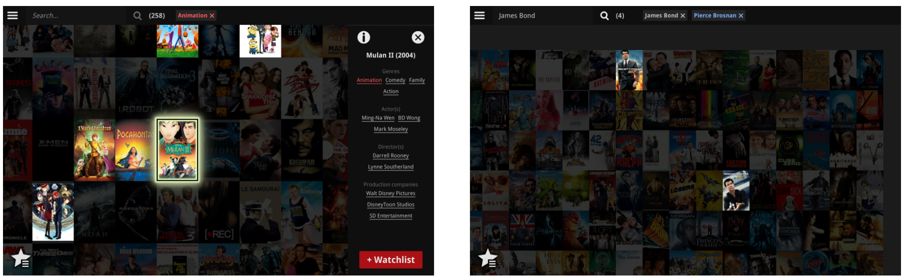

Fig. 2. The left image shows the movie wall with the genre *Animation* applied from the side panel. The right images shows multiple filters: *Pierce Brosnan* and *James Bond*, which highlights all movies starring the actor *Pierce Brosnan* and containing the keyword *James Bond*. *James Bond* is a custom filter in this case.

Arrangement of movies We tested different arrangements of the movie collection in the grid: randomly distributed and (semi-)clustered based on genre. No significant differences were found between the different arrangements, although most participants stated a preference for a clustered arrangement. Yet, one-sixth of the participants did have a strong preference for the random arrangement and some others stated they would use either based on the situation. Hence, we see value in providing both and leaving users to choose whichever they prefer. The clustered arrangement is shown in Fig. 3.

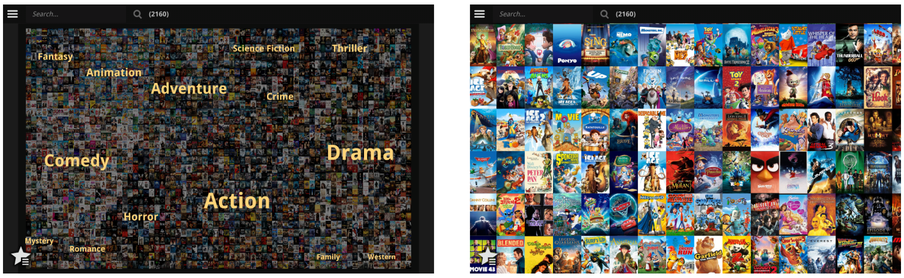

Fig. 3. Arrangement of movies clustered by genre. The left image shows it completely zoomed out, with labels that indicate the position and size of each cluster. The right image shows the same arrangement, zoomed in on the *Animation* cluster.

3 Conclusion and Future Work

We presented an interface concept that facilitates exploratory browsing of large movie collections, aimed at complementing existing, state-of-the-art movie browsing services. So far, we focussed on usage of the interface on mobile devices, tablets to be precise. For future work, it would be interesting to expand this research to different devices, e.g. smartphones and computers. Furthermore, while we designed and tested the interface within the domain of movie browsing, we are positive this interface can be applied to other domains as well, as it can essentially provide a starting point for the exploration of any (largely) visual collection.

References

[1] Wolfgang Hürst et al. "Size matters! how thumbnail number, size, and motion influence mobile video retrieval". In: *Advances in Multimedia Modeling* (2011), pp. 230–240.

[2] Antonio Torralba, Rob Fergus, and William T Freeman. "80 million tiny images: A large data set for nonparametric object and scene recognition". In: *IEEE transactions on pattern analysis and machine intelligence* 30.11 (2008), pp. 1958–1970.

Chapter 3

Annotated appendix

1 Literature Study: Facilitating and Evaluating Exploratory Movie Browsing and Related Fields

1.1 Introduction

In this literature study, the state of the art and research opportunities in the field and related fields of exploratory movie browsing are identified. This will support and guide the thesis project which aims to (1) create an interface which facilitates the exploratory search of large movie collections and (2) scientifically evaluate this interface. This interface should have the following characteristics:

- It facilitates the exploratory browsing[1] of movies.
- It provides an enjoyable browsing experience to the user.
- Its core is a storyboard with movie posters.
- It is designed for use on mobile touch-based devices (especially tablets).

As there is little literature specifically targeting exploratory movie browsing, two closely related fields will also be reviewed: image and video browsing. Even though this project focuses on three main criteria (exploratory search, movie browsing, use on mobile devices), not all literature discussed in this survey will meet all of these criteria. However, all literature included contains relevant information for this project.

Section 1.2 will discuss movie browsing interfaces. As image and video browsing are closely related to this topic, both will be discussed in section 1.3. Section 1.4 proposes evaluation methods for an exploratory movie browsing system. Section 1.5 will conclude the literature study.

1.2 Exploratory Movie Browsing

Popular movie streaming services (e.g. Netflix[2]) often offer its users the possibility to search for movies using keywords or pick a movie from the recommended section. Such recommendation systems are usually the focus of improvement of these services; Netflix even held a competition to

[1] Exploratory browsing is the opposite of directed search, as for the latter a clear search target exists. We define exploratory browsing here as a search activity where the user has no or only a vague idea of a target. This can either be caused by the lack of knowledge of the search domain or simply by a desire to explore.

[2] https://www.netflix.com/

improve theirs[3], the winning algorithms described in [16, 42]. While this means recommendation systems are often quite sophisticated, they are still far from perfect. And if the user wants to browse beyond the recommended movies, there is generally no convenient way to do so. Many movie search interfaces are not focused on exploratory search and thus do not facilitate it very well. So even though there is a lot of research in the field of recommendation systems, there is little focus on providing the means for exploratory browsing. And this while exploratory browsing can be a very good tool when the search goal is vague. Moreover, it can provide a fun and entertaining browsing experience. One of the few interfaces proposed for exploratory search of movies is MovieClouds[21, 8]. MovieClouds allows its users to browse through a cloud of words, topics (especially moods and emotions) extracted mostly from the audio and subtitles of the movies. Participants in a user study thought it was a interesting and fresh visualization of movies, but the interface was a bit overwhelming. Another approach to exploratory movie browsing is shown in [19], which presents a graph-based interface called NEMP. The user is presented with a few popular movies at the start (using movie posters) and can choose a movie to its liking, whereas the interface consequently shows similar movies to the chosen movie, based on metadata (genre, actors, etc.). This process can be continued for as long as desired. The interface was compared to the TMDb web interface[4] and for both interfaces participants were asked to add movies worth watching to their personal watchlist. There was no significant difference in number of movies added to the watchlists with each interface. Users indicated though that interesting links could be found more easily using the NEMP interface.

This outlines how little research has been done in the field of exploratory movie browsing. Furthermore, both interfaces shown could not really prove that their performance was better than any other movie browsing interface already existing, perhaps not caused by the lack of good initiatives but by the means of evaluation. This shows the importance of a method to properly assess the performance of such systems, a subject touched upon in section 1.4.

1.3 Image & Video Browsing

While not much literature focuses on movie browsing interfaces, plenty can be found on image and video browsing. Both fields are closely related to movie browsing, as of course a movie is a specific type of video and its movie poster is a specific type of image. Moreover, videos and movies are a medium often represented in interfaces by still images (key-frames). The difference should be noted as well though, namely the fact that each movie, unlike much other video content, already has a representing thumbnail available in the form of a movie poster. And other than most image and video content, there is much more metadata available for a movie (poster). Still, much can be taken from image and video browsing research.

Many (commercial) image and video browsing systems we use today, feature a 2D storyboard interface, where the user can scroll through thumbnails presented in a grid-like structure, e.g. Google image search and many standard mobile gallery applications. However, interfaces like these have their limitations, the first one being that the user can only see a relatively small part of the images or videos in the collection at all times. This makes it difficult to get a good overview of the entire collection and its structure. Other than that, it makes it hard or even impossible to compare items far apart. Furthermore, to browse the entire collection, the user has to go through tedious scrolling actions. When not finding the desired image, the user has to go back and forth or start again, which can be frustrating. Lastly, there is no way that the user can find an item located at the end just as fast as an image located at the start of the collection[36].

Various interfaces and improvements have been proposed the past years. Section 1.3.1 shows research on similarity-based structured interfaces. Section 1.3.2 discusses three main different methods to visualize an image or video collection. Section 1.3.3 discusses 3D browsing interfaces. Lastly, section 1.3.4 contains research on the use of (multiple) small thumbnails.

[3] The Netflix Prize was an open competition where competitors had to improve the prediction accuracy of how much someone would like a movie. See http://www.netflixprize.com/ and [3]

[4] The Movie Database website: https://www.themoviedb.org/

1.3.1 Similarity-Based Clustering

This section contains different methods of ordering (large) collections of images or videos. We will not go in-depth about the technical background (apart from its consequences for performance) but mainly look at the high level concepts and their advantages and drawbacks.

Visual similarity
Many image browsing interfaces use some sort of similarity measure to determine the ordering of images. In [28] the usefulness of similarity-based clustering is evaluated. A grid with randomly distributed images is compared to an interface with images sorted by visual similarity (based on low-level visual features), showing that users were able to locate specific images faster using the latter. These findings are confirmed in [18], where an interface based on visual similarity is proposed to effectively browse image results on the web. In [27] however, participants were slower with a visual similarity-based arrangement, indicating that a random arrangement could work better in some cases. An explanation for this is that using an arrangement based on visual similarity, similar images sometimes seem to merge, making them less eye-catching than in a random arrangement. However, there was a preference for the similarity-based arrangement though not overwhelming; participants thought it was more enjoyable to use and thought this made it easier to find the photos they wanted. A drawback of clustering based on image features is that these methods are often expensive[30], which could make them infeasible for use on less powerful, mobile devices.

A specific type of visual similarity-based sorting is sorting based on color, used in several browsing interfaces [1, 29, 33], as it was shown that this can improve known item search of images [31]. An advantage of color sorting is that it is intuitive and fast, the latter meaning that it is applicable for mobile devices as well.

Caption-based similarity
Other than based on visual features, image and video content can also be sorted on captions. Images and videos often have some semantic context which is not captured in visual features, so for some content it might not make sense to sort on visual features. Movie posters can be an example of this, as shown in [15]. In [27] a caption-based arrangement was shown, which was actually slightly preferred over their visual-based sorting. It was not tested against a random arrangement though. While caption-based similarity arrangements can be really powerful, they can only be as good as their captions. And an issue with this being that many images and videos lack these (good) captions. Furthermore, labels might be needed for the user to understand the structure of a collection. The need to label clusters for clarity was also mentioned in [5, 40].

We can conclude that similarity-based visualization is a powerful approach as a user can quickly assess whether a cluster in its entirety is of interest to her or not by only looking at a few items [19]. Caption or context-based similarity seems the most potent approach for arranging movies, as it might not make sense to sort movie posters on visual similarity[15]. Furthermore, movies often have a lot of metadata available that can be used for clustering (genres, actors, director, etc.), as done in [19]. We have seen that a potential drawback of visual-similarity sorting is that certain items can be overlooked when placed in a similarity-based arrangement as opposed to a random one. Whether this characteristic transfers to caption-based similarity arrangements as well cannot be assumed without further investigation. Randomness could lead to more unexpected discoveries, but might not allow the user to direct her search well enough. All in all, the best way to arrange movies is not at all obvious and provides an interesting research topic.

1.3.2 Visualization of Image and Video Collections

Image and video browsing interfaces using similarity-based sorting can generally be divided into three categories: distance-based, hierarchical and graph-based. This section discusses the visualization methods, in order to assess their potential suitability for an exploratory movie browsing interface.

Distance-based visualization
While interfaces often use an uniform mapping of thumbnails, there are also those that position thumbnails based on their differences, showing similarity using spatial mapping, e.g. [28, 26, 18, 39, 40, 9]. These often use methods like principal component analysis (PCA) or multidimensional scaling (MDS) to calculate Euclidean distances between images. This also touches upon the first drawback of these methods: they can be computationally expensive, especially when calculating global maps. The second potential drawback is that, even though these methods provide a very clear clustering of the collection, it also introduces overlap of images. While not being the focus of the research, in [28] every single test subject disliked the overlap in the similarity-based interface. According to them, overlap made it hard to see the edges of an image and sometimes occluded important details, making images less recognizable. In [18] users could adjust the overlapping rate, being able to choose between preservation of the spatial differences or less overlap. Unfortunately, they do not show any in-depth data about the use of this feature, other than that users appreciated being able to adjust the level of overlap. A user study of [39] showed a slight user preference for a uniform (grid) layout over a non-uniform layout, but there were no clear differences found between the two. However, for touch-based devices overlap could introduce another issue as selecting on a touchscreen device is less precise than with a mouse, and selecting an image in an overlapping cluster could prove to be much harder[9].

Hierarchical visualization
With large collections, it is difficult to show all the content initially. Putting everything in a standard 2D scrollable storyboard will force users to do a lot of tedious (horizontal) scrolling in order to find what they are looking for. In contrast to this, hierarchical interfaces only show a part of the content, ideally being representative for the entire collection. Typically the user can then zoom in on a region of interest, revealing more, similar images or video (key-frames). This hierarchical nature makes such interfaces suitable for storing large collections. For example, the hierarchical globe interface of [29] can theoretically store over 23 million images with just 3 hierarchical levels. DynamicMaps[14], a hierarchical infinite 2D grid interface, can also handle millions of images. But it should be noted that this is also thanks to its local map generation, overcoming performance issues with commonly used (global) clustering methods. A user study comparing DynamicMaps to simple Relevance Feedback interface, indicated that it supported exploratory search well, and was more fun to use. For video browsing hierarchical structuring can also be very useful. An example of this is shown in [38], featuring an interface with 3D rings with key-frames of a video. The user can inspect video segments in as much detail as needed, opening a new ring for each hierarchical level.

While hierarchical interfaces have its advantages, it can be hard to for users to decide which region contains what they are looking for. A user study of [29] showed that with a Known Item Search task, users had a hard time finding the right image if it was not on the top level.[1]. Furthermore, it can be hard to find images with less common concepts, as they can be hidden away in another area[14].

Graph-based visualization
Graph or link-based visualizations enable the user to browse a collection step by step based on similarity. The user can typically choose a thumbnail in the current visualization and new similar, linked content will appear instead of the old ones. The berry-picking approaches in [20, 19] use this concept combined with distance-based visualization. They calculate local maps of the neighborhood and align them to get the impression of panning through a global map, each time selecting a new image (or movie). They argue that these local maps are more accurate than a global map and less expensive. The Rotor-[7], Cross- and ForkBrowser[6] for browsing video key-frames also use a link-based approach. They use threads, lists of shots sorted on some feature, e.g. query result, visual similarity, semantic similarity or time. The interface visualizes these threads as different directions the user can go. When a new shot is chosen from the available threads, new threads linked to that shot appear. A hierarchical graph approach is shown in [2], for exploratory image and video search. The interface shows the previous image, current image and all images three levels deeper in the graph. The further in the graph, the smaller the image. To accompany this, they also have a grid-like visualization where all the graphs are sorted around

the current image.

A drawback of graph-based visualizations is that the user never really has an overview of the entire collection. But one could argue that this is hard to accomplish with any interface when the collection gets really large, which is the case with many image and video collections on the web these days. Another potential issue of graphs is that it is very possible to go around in circles, which is usually not desirable. Also, going back or finding the right link could be tedious if the user wants something entirely different than her current location in the graph.

All three have their own advantages and drawbacks. Distance-based visualizations create clear clusters, enhancing what was mentioned earlier about being able to judge an entire cluster on its relevance by only seeing a few items. Drawbacks are however that global distance calculations can be expensive and the overlap these it creates is potentially undesirable. It's also hard to visualize a large collection with this visualization, but it can be combined with the others, for example shown in [20, 19]. Graph-based visualizations can work well for large collections and is relatively inexpensive if local maps are used[19]. It's a good way to keep refining a query when on the right path. However, it is hard to get an overview of the entire collection and it can be tedious to get on the right path when currently in a wrong part of the graph. This is less the case with hierarchical visualizations, as often the user can simply zoom out and go to a different region of interest. Because of this zoomable nature it could also prove to be very suitable for touch-based devices as zooming is often a very easy and intuitive interaction on these. It can be a difficult task though to represent the collection well on the top-level of the hierarchical structure. Rare concept might be hidden away somewhere unexpected and hard to find. We have already seen the graph-based approach in movie browsing[19], but the hierarchical approach only in image browsing. Yet it showed promising results in an exploratory context, e.g. in [14], so whether this works for movie browsing on a mobile device as well is another research opportunity.

1.3.3 3D Interfaces

In addition to 2D interfaces, many 3D interfaces for image and video browsing have been proposed the past years. Making use of a 3D visualization space can result in better use of available screen space and can thus show more images or key frames at a glance[36]. The MediaMetro[4] and mTable[5] interfaces use a 3D city metaphor for browsing image and video content. In [38], a 3D carousel is proposed for browsing videos, featuring hierarchical levels to view video content in as much detail as desired. An interface with several 3D layouts is shown in [22], so users can choose the layout of their preference for their current task. MINI[9] is a mobile 3D image browser for multi-dimensional data. The use of multiple planes allows them to sort data on different criteria. A video browsing interface for tablets is presented in [11], featuring a 3D filmstrip to browse through key-frames. Unfortunately, these papers contain little to no evaluation or comparison to the standard 2D storyboard interface. Two 3D interfaces that have been evaluated more extensively are the Ring and Globe interfaces.

The Ring interface started out as a vertical cylindrical 3D storyboard [37, 35]. The advantage of this interface over the standard 2D storyboard is that it is able to show a larger part of the collection, with different levels of detail. Low detail for the images in the back, high detail for the images in focus. This makes it easier to compare or navigate to images in different locations. A user study comparing the 3D interface to a standard 2D storyboard interface showed only a small, non significant improvement in performance though. The Ring was later improved upon, by using color sorting, making it horizontal and using bended images in stead of flat ones[34]. Even though the latter resulted in more distorted images, there was less occlusion and got rid off gaps in the cylinder. An advantage of the horizontal Ring is that it gives a good view of the images in the back, without too much distortion Other than that, the user study indicated that it was more natural looking than the vertical version. No significant difference in performance between the two was shown, but most users preferred the horizontal ring. In [32] the Ring was compared to a standard 2D storyboard featuring the same color sorting. The Ring was 12.7%

faster with Known Item Search tasks. It has to be noted though, that the trial time was very dependent on the location of the target image in the interface. E.g., target images located on the sides of the ring as seen from the initial state (where the images were almost occluded) took significantly longer to find.

The Globe interface was introduced by [29]. The interface consists of a color-sorted globe with hierarchical levels. The authors argue that a globe provides an intuitive interface, as a globe is a familiar concept because of its direct analogy to the earth. In [1] two improved versions of the Globe are presented and tested. This resulted in an improved version which was no longer hierarchical and used a different color-sorting method with special regions for very dark and very bright images (H-Globe). The hierarchical levels didn't work very well for Known Item Search as images that were not located on the top level turned out to be hard to find. In a user test with 24 participants, a set of 350 images and a Known Item Search task, the H-Globe outperformed the standard 2D storyboard with 23%. It also scored better ratings for workload (NASA-TLX).

We've already seen that the Ring and Globe both outperform a standard 2D interface on a PC. In [10, 36], the Ring[33], Globe[1] (H-Globe) and two other interfaces are compared in several tests, on several devices. The Ring and Globe performed very similar when used on a tablet. In a second test, both were also compared to the Grid and ZoomableGrid on a small mobile device. The Grid is a standard 2D storyboard interface. ZoomableGrid is an interface where each image is visible initially in a grid-like structure and can be zoomed in on to see images in higher detail. Each interface applied the same color sorting. User tests took place with different set sizes (100, 200, 300, 400). The Grid performed worse the higher the set size, which can be easily explained by the extra scrolling needed with higher set sizes. The other 3 interfaces did not have this issue. The Ring and Globe once again performed similar and had similar flaws. One problem with the 3D interfaces is that images at the sides of the 3D shapes at the initial view were almost entirely occluded. ZoomableGrid performed best, outperforming the others in most set sizes. The authors see ZoomableGrid as the most suitable for small mobile devices and recommend looking into combinations of the ZoomableGrid and the 3D interfaces for future work.

Altogether, 3D interfaces definitely have potential as their main advantage is that they are able to show more images at a glance, in different levels of detail. The Ring and Globe interfaces were able to outperform standard 2D storyboards on PC and tablets. The characteristics of the specific 3D shape used has to be taken into account though, as the images at the sides are easily occluded. However, it is hard to say whether they provide advantages for exploratory search, as none of the 3D interfaces above have evaluated this type of browsing. Also, it need to be noted that the interfaces that have been evaluated, have often done so using small set sizes of a few hundred images or videos.

1.3.4 Small Thumbnails

Even though many interfaces shy away from using small thumbnails, it has been shown that humans are very capable of recognizing small images, especially when it concerns moving images [12]. These findings partly inspired the interface concept in this project, being confident that users can deal with multiple, relatively small images. In [41] it was shown that humans were able to accurately recognize more than 80% of 32×32 pixels color images on a PC, outperforming computer vision algorithms. Similar results were shown in [12], when testing different thumbnail sizes on a mobile phone. Especially for moving thumbnails the results were surprising, at 30 pixels the error rate was lower than 20%. As design guidelines, they suggest using at least 90 pixels for static thumbnails and 70 pixels for moving thumbnails. In [13], it was tested whether these values still hold when multiple thumbnails (5-9) are shown instead of one, isolated image. This turned out not to be the case, however, their results indicate that users are still very capable of dealing with multiple small thumbnails. A user study comparing several browsing interfaces[35] also found that when using an interfaces with images of 55×32 (4 times smaller than images in the 2D storyboard they compared it to) there was no significant difference in performance. It has to be noted though, that it performed significantly worse in workload and was not very liked by the participants.

All in all, this research shows that humans are definitely able to handle multiple small thumbnails, not even at the cost of significant performance. However, the question remains if they would like to do so. Earlier work seems to indicate that user do not like it if they can only browse using small images[35], but it was fine in a situation where also more detailed images were available, as was the case with the 3D interface featured in the same paper.

1.4 Evaluation

For the evaluation, it is important to keep in mind that the envisioned interface for this project focuses on the following:

- Facilitating exploratory browsing of movies.
- Providing an enjoyable browsing experience.

Evaluating an exploratory search interface is quite different from evaluating a system focused on directed search. For the latter, efficiency is often paramount and thus what is being evaluated. This can be done measuring objective data like trial times and error rates. For exploratory search interfaces however, this is not so easily done as there is no clear, desired outcome[24]. Moreover, in a direct search scenario the user often desires a fast result, while in an exploratory search scenario this is not always the case, as the user might just enjoy browsing around for a while, so objective metrics like time to find a target do not indicate performance of such a system.

Often researchers use their own, non-standardized questionnaires with statements that participants rate on a Likert scale, to get an idea of user satisfaction, e.g. in [19, 14]. However, this is quite limited and imprecise[24] and it can be hard to draw real conclusions from those about user experience. Many also use the NASA-TLX to measure subjective workload. While less workload can indicate a better browsing experience, it's not enough to conclude this.

To evaluate whether a system provides an enjoyable browsing experience, one should assess the user experience or engagement of such a system. Especially for exploratory search, as user engagement is a strong indicator of the performance of an exploratory search system.[43]. User engagement can be measured with the following approaches: self-reported engagement, cognitive engagement or online behaviour metrics.[17] The last one requires vast amounts of data which is hard to obtain within the scope of this project. Although cognitive engagement allows working with objective data, the interpretation of such data includes a whole other research domain. The remaining option is self-reported engagement, which relies on user subjectivity. It can be used in both lab settings and online studies and seems the most feasible for this project. A method to do so is the User Engagement Survey (UES)[25, 24], using six dimensions of engagement: Aesthetics (AE), Focused Attention (FA), Felt Involvement (FI), Perceived Usability (PUs), Novelty (NO) and Endurability (EN). It was initially validated using e-commerce environments, but later its generalizability in the field of exploratory search was examined[24] and the model was also adapted for use with mobile devices[23].

1.5 Conclusion

In this literature study we have examined exploratory movie browsing and its related fields: image and video browsing. It has shown that there is little research in the area of exploratory movie browsing and even less in the evaluation thereof, leaving many research opportunities. The field of exploratory searching itself is also still developing as (commercial) browsing interfaces still mostly rely on keyword search. This we've also seen when examining video and image browsing interfaces; many are focused on directed search and are evaluated accordingly, using objective measures like trial times and error rates. For the exploratory systems discussed here, most evaluations use self-engineered questionnaires, which are limited and less reliable when concluding anything about user engagement. To really measure the multi-dimensional concept of user engagement, better methods are necessary, e.g. the User Engagement Survey[25].

Literature concerning image and video browsing has shown us valuable lessons about sorting and visualization techniques, which should be kept in mind during the creation of the interface envisioned in this project. However, as the fields still have their differences and many interfaces were focused on directed search, we cannot simply transfer each finding. It has been shown that similarity-based clustering can assist image browsing, especially using visual similarity. For movie browsing however, the question remains whether clustering on metadata (e.g. genres) will assist this activity or will narrow its users vision, preventing them from doing serendipitous discoveries. The same field has shown us that structuring content using (Euclidean) distances, graphs and hierarchies can ease and improve the browsing process, the last two especially with large collections. However, we have only seen the graph-based visualization in the field of movie browsing[19]. Could a hierarchical visualization, e.g. as shown in [14] prove to be useful for an exploratory browsing interface? Furthermore, 3D interfaces have been discussed but not yet tested with exploratory search scenarios. Could 3D interfaces have an advantage over 2D interfaces when used for exploratory browsing? These are all questions that to the best of our knowledge have not yet been answered, and thus provide interesting research opportunities.

2 Pilot Study

A pilot study was conducted in the earlier stages of this project to verify the potential and feasibility of the *MovieWall* interface concept. The main results of this user study are reported in chapter 1, section 3.4. This section contains additional implementation details and a more elaborate overview of the data obtained during the user study.

2.1 Implementation

The prototype application was developed in Unity[5]. At early stages it was already decided that the application was going to be very visual and possibly contain visual animations, which is relatively easy to achieve with a game engine. Furthermore, with Unity the application can easily be ported to several platforms, including Android.

The application is designed for use on touch-based tablets, in landscape orientation. The resolution of the test device (Samsung Galaxy Tab A 10.1), 1920 × 1200 pixels, was used as a reference resolution.

The interface features a large grid of 8640 movie posters in total. A low resolution image (62 × 93 pixels) of each movie poster is stored locally, to decrease the number of web requests needed during runtime. Because this low resolution becomes very apparent when zooming in beyond a certain image size, high resolution images are loaded during runtime when the movie posters on screen become larger than 124 × 186 pixels. These high resolution images are 185 × 277 pixels. These high resolution images are loaded in memory and disposed of after a while of not being used to limit memory usage. Only the low resolution images are stored locally to limit the application size.

Both the poster images and movie data used were obtained from The Movie Database, a community managed movie and TV database (*www.themoviedb.org*). This database was selected both for its accessible and easy to use API, extensive documentation and the comprehensive amount of metadata per movie. The one drawback of using their API is the request limit, which allows no more than 40 requests every 10 seconds. However, since image requests do not fall under this request limit, it should not be an issue, as it is very unlikely that a user desires and succeeds in selecting 40 movies within 10 seconds.

2.2 Data

Table 3.1 show the ratings given by each participant for each statement in the questionnaire. Explanations of the ratings given are summarized in Table 3.2. Tables 3.3-3.7 contain additional data from the interview. Figures 3.1 and 3.2 show graphs of the zoom levels used during the experiment by each participant. We use **P1-P10** to refer to each individual participant. In the tables that contain the interview data, we use the **#**-sign to indicate how many participants gave a similar answer or explanation.

Table 3.1: How each participant rated the four statements given, after using the interface. Each statement was rated on a 5-point Likert scale.

Statement	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10
I was overwhelmed by the interface	4	3	3	2	4	4	4	4	4	5
I thought the interface was easy to use	5	5	5	4	5	5	5	5	4	5
I thought interacting with the interface was intuitive	4	5	5	5	5	4	5	4	3	5
I thought browsing the interface was enjoyable	3	3	3	4	5	4	4	4	2	2

[5]Unity is a cross-platform game engine, free to use for the development of non-commercial applications.

Table 3.2: Interview questions and answers according to the ratings given to each statement. The exact ratings given for each statement by each participant are shown in Table 3.1.

Statement: *I was overwhelmed by the interface.*				
Rating	**#**	**Question(s)**	**Answers**	**#**
1-2	1	You didn't think the interface was overwhelming. Why do you think you were able to handle it well?	It was clear that you could zoom in, making it less overwhelming	1
		Why did you think the interface was (slightly) overwhelming?	There are a lot of movies on your screen initially No clear structure I didn't know where to start	8 3 3
3-5	9	Do you see this as a negative aspect of the interface?	Initially, yes, I didn't know where to start After a while, yes, because I wanted more structure. Yes, if it would always be like this, for once it's fine Yes, it is too much No, but perhaps it would be good to show that you can zoom in initially No, it was a lot initially, but after zooming it was fine No, you see how much choice there is	2 1 1 1 1 2 1
		Do you have any ideas on what could/should be done to make it less overwhelming?	Zoom in initially Filter functionality, either during or before browsing Cluster the movies in subcategories Add indication that it is possible to zoom Use different sized images	4 3 2 1 1

I thought the interface was easy to use.				
Rating	**#**	**Question(s)**	**Answers**	**#**
4-5	10	You found the interface quite easy to use. Any comments on that?	Controls are simple and straightforward Very intuitive to use Easy to use if you know how a tablet works Functionality was clear Works fluently It worked like I expected it to	3 3 2 1 1 1

I thought interacting with the interface was intuitive.				
Rating	**#**	**Question(s)**	**Answers**	**#**
1-3	1	Why didn't you think the interface was intuitive to use?	Initially it wasn't clear to me what I should do.	1
		How would you make it more intuitive to use?	Adding a tutorial Automatically zooming the grid to some point	1 1
4-5	9	Why did you think the interface was intuitive?	It uses standard gestures for zooming and scrolling It worked like I expected it to	5 4

I thought browsing the interface was enjoyable.				
Rating	**#**	**Question(s)**	**Answers**	**#**
1-3	5	Why didn't you think browsing the interface was enjoyable?	It was hard to look for something specific You can't see similar movies of an interesting movie There were too many movies I felt like I was just aimlessly scrolling through the application There were a lot of uninteresting movies I didn't know where to start	4 2 2 1 1 1
		Can you think of anything that could/should be done to make it more fun or enjoyable?	More structure, i.e. sorting on genre Adding filters Being able to save a movie or go back to a previous movie No, but it's practical right now and that is fine	2 1 1 1
4-5	5	Why did you think browsing the interface was enjoyable?	I liked just casually browsing through the movies It brought me to new/unexpected movies I saw more than I normally would It helps to find a lot of movies	3 3 1 1

Table 3.3: In the interview participants were asked to explain certain observed interaction behaviour (zooming, side panel usage). The explanations given are shown in this table.

Interaction	Option	#	Explanation	#
Zooming	Not used very often	8	I reached a level where I could view the posters well and kept using that	4
			It didn't feel useful to use zooming to get to another part in the grid because of the random arrangement	2
			Just scrolling was fine to discover new content	2
	Used regularly	2	Zooming out allowed me to go through the movies faster	1
			I used it to navigate to a different region	1
			I zoomed out a little more to look at the movies globally, if something seemed interesting I zoomed in on that	1
Panel	Often closed	9	I wanted to have the space to look at other movies	6
			I only used it to check the details	6
			I can't handle that much on my screen	1
			I judged the movies by posters, so I didn't need the panel	1
	Often opened	1	To compare the selected movie to others	1

Table 3.4: Participants' answers to the question: *"Do you think the random arrangement* [of the grid] *was good, or would you have an idea for a different arrangement that might work better?"*

Opinion	#
It would be nice (and less overwhelming) if there was more structure	3
It would be nice to be able to filter out categories I don't like	3
In the beginning it's fine, after a while I want to be able to direct my search more	2
It can be fun, but it is not something I would always want	1
I expected it to be sorted on genre and I think this would be nice, so the search can be a bit more directed	1
It's nice that you're more stimulated to look for different genres	1
I didn't miss the lack of structure	1

Table 3.5: Participants' answers to the question: *"Do you think this interface could be an alternative or complement to existing movie browsing interface (e.g. of streaming services)?"*

Answer	#	Explanation	#
Either	1	Yes, I was really able to find something new with this interface	1
Alternative	1	You have a bigger overview of what there is on offer and you see more possibilities	1
		It would be a nice addition when you don't really know what to watch	1
Complement	7	But only with additional ways to direct your search	1
		Yes, but I wouldn't use it	1
Neither	1	I like the way it works, but right now there is too much input and no way to direct my search	1

Table 3.6: Participants' answers to the question: *"In what situations would you rather use this than existing movie browsing interfaces?"*

Answer	#
When I don't have a real preference	4
When I don't know what to watch	4
When picking a movie with multiple people (friends, family)	3
If I wanted to watch something less known	1

Table 3.7: All the suggestions done by the participants during the interview.

Suggestion	#
(An option to) Sort/cluster the grid	5
The option to filter out content you don't like	4
Search function	3
Show related movies, somehow	3
It would be cool to make the metadata in the panel interactive	2
Infinitely scrollable grid (or something else that would fix the movies at the border not being in focus)	2
Spacing between the posters	1

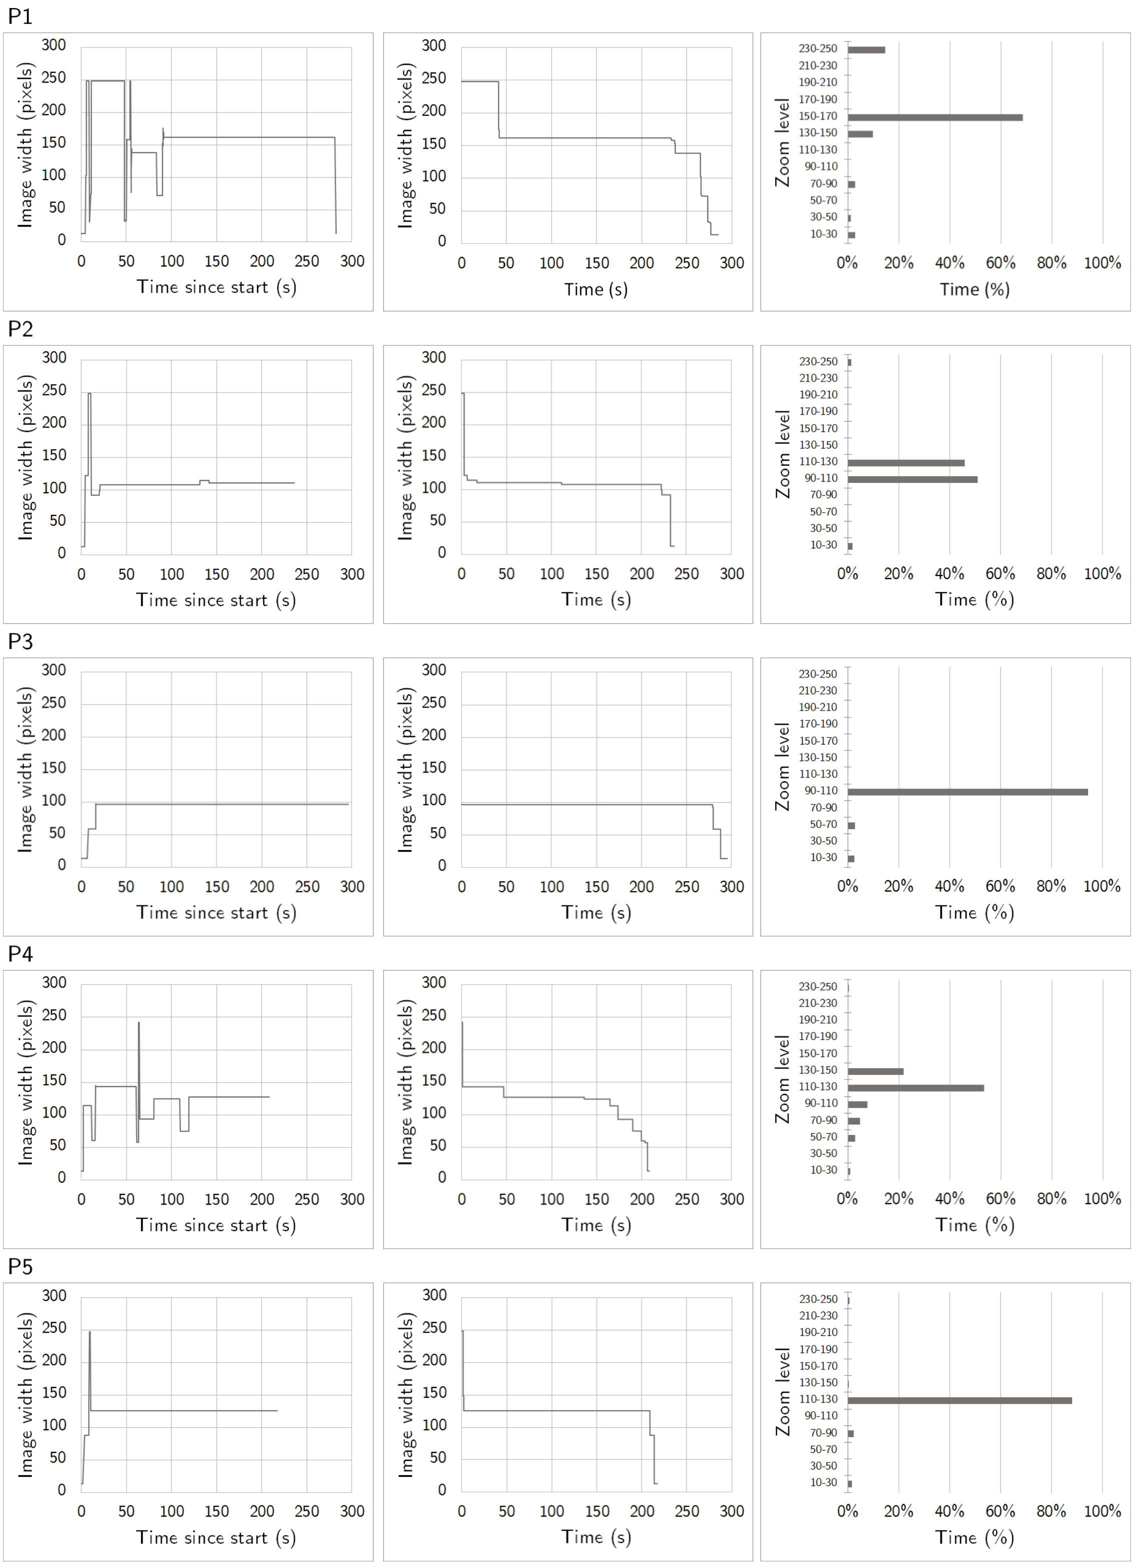

Figure 3.1: Visualization of the zoom levels used by participant P1-P5. (1) The first graph shows the zoom level over time. (2) The second graph contains the same information, but sorted by zoom level. (3) The third graph shows the percentage of time on each zoom level, where the zoom levels are grouped in bins.

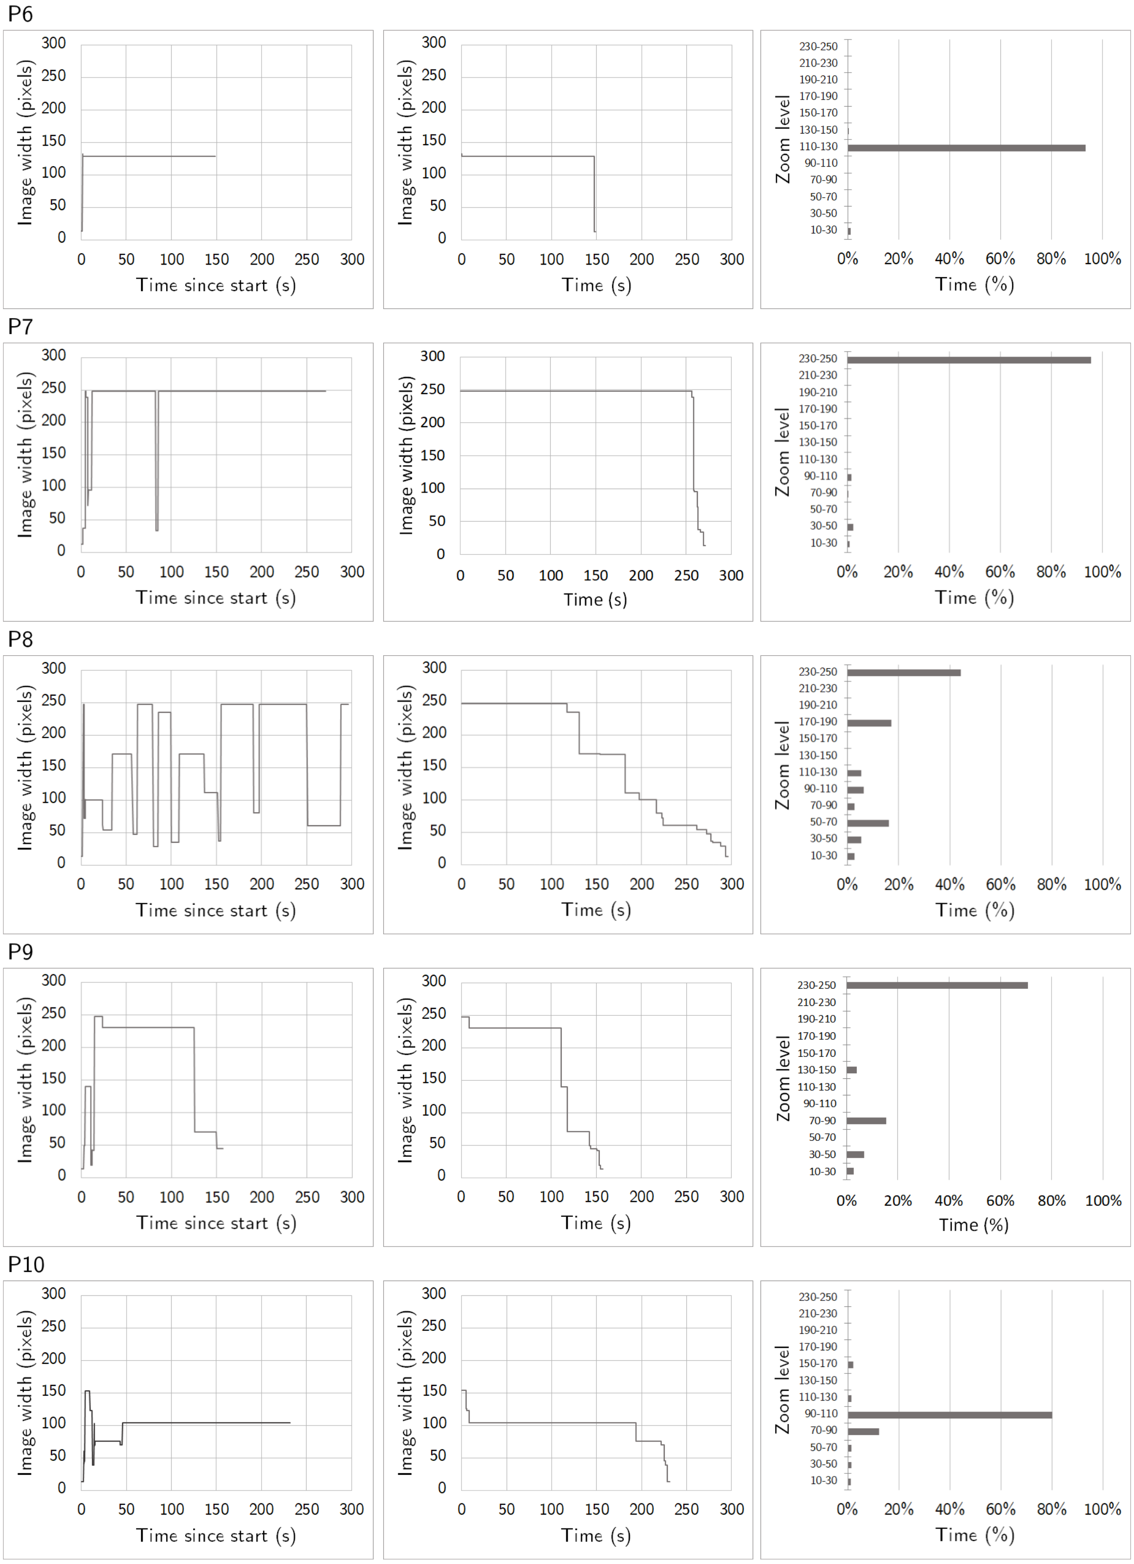

Figure 3.2: Visualization of the zoom levels used by participant P6-P10. (1) The first graph shows the zoom level over time. (2) The second graph contains the same information, but sorted by zoom level. (3) The third graph shows the percentage of time on each zoom level, where the zoom levels are grouped in bins.

3 Detailed User Study

Based on the findings in the pilot study, the interface was improved and a second, more detailed and formal user study was conducted. The main results of this user study are reported in Chapter 1, section 4.4. This section contains additional implementation details and a more elaborate overview of the data obtained during the user study.

3.1 Implementation

The main differences of the improved interface are the filters and the additional (semi-)clustered arrangements. For the filters, a local database had to be created, as described in Section 3.1.1. Furthermore, two clustered arrangements were added for which a clustering algorithm had to be created, as described in Section 3.1.2.

3.1.1 Local database

For the filter functionality, it had to be possible retrieve specific data of each movie in the grid at once. For this it was extremely inconvenient to make web requests to the API at runtime, because of the request limit mentioned in section 2. Hence, we created a local SQLite database with any necessary data of the 2160 most popular movies in The Movie Database.

3.1.2 Clustering

For the clustered arrangements we decided to cluster the movies based on genre. This seemed the most intuitive sorting and provided an acceptable number of clusters as there are 19 different genres in The Movie Database. We decided to only put each movie in one cluster, to keep the movie collections of each different arrangement comparable. The placement of clusters relative to each other is random. We considered placing highly correlated clusters close to each other, but this is not a trivial task because of the numerous correlations between genres. Furthermore, using the zooming functionality, users can easily swap between genres. Hence, we did not think it would significantly influence the outcome of the user study if we used a more advanced clustering algorithm.

The clusters were created using an external tool. The tool works as follows:

- The first (most prominent) genre of each movie is extracted from the database.
- For each cluster, starting with the genre with the most movies, the “best” location is determined. This is done selecting a random index in the grid, checking the x closest indices around it, where x is *the number of movies in the cluster* and counting how many of those indices are already in use by other cluster(s) (overlap). This repeated a 100 times or until the overlap is below a certain threshold. In the end the best option is selected, and the cluster is added to the grid. Figure 3.3 (left) shows a possible outcome.
- Because this outcome is often far from optimal as can be seen in the image, it can be optimized. In one optimization iteration, it is checked for each index whether it is beneficial to swap it with another index in the grid. If so, it is swapped with the most beneficial option. A swap between two indices is beneficial if the sum of their distances to each respective cluster center becomes smaller than what it currently is.
- Although this simple approach gets stuck in a local maxima after some iterations, the outcome is often satisfactory. An example of an optimized solution is shown in Figure 3.3 (right).
- If desired, the outcome can be further adjusted manually and saved to a .json file.

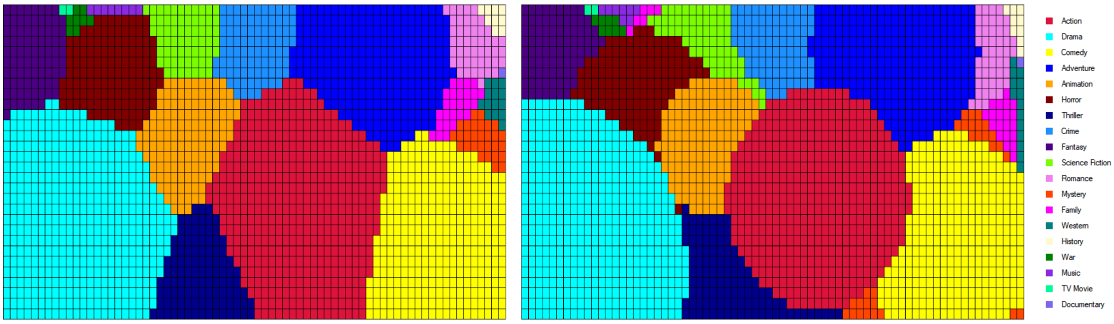

Figure 3.3: Possible distribution of clusters, calculated and visualized by a separate tool. Left shows an initial, non-optimized solution. Right shows an optimized solution, based on distance to the cluster centers.

In the application, the file with cluster data is read and all movies are distributed accordingly. Furthermore, within a cluster, the movies are sorted on popularity, with the most popular movie in the center of the cluster. Additionally, labels with genre appear when zoomed out far enough, as shown in in Figure 3.4. The label size scales with the cluster size and only appear for clusters with a minimum of 10 movies.

3.1.3 Semi-clustered arrangement

For the semi-clustered arrangement the same cluster data was used as with the strictly clustered arrangement. However, a certain percentage in the border area of each cluster is randomly swapped with another movie in the grid. We chose to only swap movies in the border, as we expected it would be strange for users to see a randomized movie in the middle of a cluster. In this case, the outer 70% of each cluster was treated at randomizable area, and of this area, 50% of the movies were randomly swapped with other movies. Additionally, within a cluster, 25% of the movies were randomly swapped, creating randomness within the popularity sorting. These percentages were chosen as they provided a desirable outcome. A possible result of this randomization is shown in Figure 3.4.

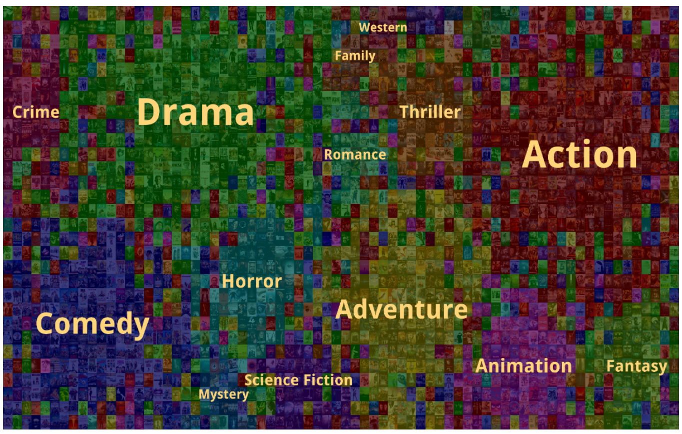

Figure 3.4: A visualization of the semi-clustered arrangement.

3.2 Data

The tables below provide additional log, questionnaire and interview data. Table 3.8 shows the interaction with the movie collection of the first trial of each participant. Tables 3.9 and 3.10 provide additional data about the filter usage. Tables 3.11 and 3.12 show the results of the questionnaire, filled in by each participant after each trial. Finally, Tables 3.13, 3.14 and 3.15 contain all interview questions and answers. We refer to participant numbers when using **P**. Furthermore, we abbreviate the three arrangements; Random, Clustered and Semi-Clustered, respectively with **R**, **C** and **CR**. In the tables that contain the interview data, we use the **#**-sign to indicate how many participants gave a similar answer or explanation.

Table 3.8: The number of unique movies inspected during the first trial and number of movies in the watchlist at the end of the first trial, listed for each participant.

Arrangement during first trial	P	Number of unique movies viewed during the trial	Number of movies in watchlist at the end of the trial
R	1	14	4
	2	12	7
	7	10	3
	8	12	10
	13	15	10
	14	13	11
	19	13	11
	20	10	3
	25	6	3
	26	4	1
C	3	10	6
	4	10	8
	9	10	7
	10	14	3
	15	10	7
	16	6	4
	21	6	3
	22	5	4
	27	12	7
	28	10	10
CR	5	7	4
	6	5	4
	11	6	6
	12	8	5
	17	10	8
	18	7	4
	23	9	4
	24	3	3
	29	7	6
	30	9	4

Table 3.9: Filter usage during each trial of each participant.

P	Total number of filters used during the trial			Max number of filters active during the trial		
	R	C	CR	R	C	CR
1	3	2	0	2	1	0
2	3	0	0	2	0	0
3	13	0	10	3	0	3
4	8	5	6	2	1	1
5	0	0	0	0	0	0
6	5	0	5	1	0	2
7	7	0	0	1	0	0
8	5	3	10	2	1	2
9	4	1	3	2	1	1
10	2	3	11	1	1	3
11	2	0	0	1	0	0
12	2	7	3	1	2	1
13	14	5	12	3	1	2
14	13	6	2	1	1	2
15	6	1	7	2	1	3
16	0	3	0	0	1	0
17	4	7	7	1	2	1
18	5	0	0	1	0	0
19	6	3	0	2	2	0
20	9	7	11	1	2	2
21	6	3	0	2	1	0
22	1	1	2	1	1	1
23	0	10	5	0	5	1
24	0	0	0	0	0	0
25	4	6	2	1	1	1
26	4	1	0	2	1	0
27	5	5	2	2	1	1
28	0	1	3	0	1	1
29	7	3	2	2	1	1
30	8	6	14	2	1	2

Table 3.10: Usage of different types of filters with different arrangements and in total.

	Genre	Cast	Director	Company	Total
R	79	21	0	15	115
C	10	24	14	4	52
CR	37	24	2	18	81
Total	126	69	16	37	248

Table 3.11: Intrinsic motivation of each participant to use each interface setting.

P	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	5.7	5.1	4.9	4.1	5.6	2.7	5.0	5.7	4.6	3.1	3.3	6.6	6.4	6.3	3.6
C	5.9	5.6	4.9	6.0	5.3	4.4	4.9	5.1	6.3	3.3	5.1	6.0	5.6	5.6	3.3
CR	5.6	6.0	4.4	6.0	5.0	4.6	4.6	5.3	6.4	3.6	5.3	6.1	4.9	6.0	3.4

P	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
R	4.3	3.7	4.1	5.4	5.6	5.3	3.3	4.4	4.1	3.9	4.4	3.9	4.1	4.4	5.9
C	4.7	5.7	5.1	6.1	5.7	4.9	3.3	5.4	4.3	4.6	2.4	5.4	5.0	5.1	4.0
CR	4.4	5.7	4.6	5.1	5.1	4.6	4.1	5.9	4.6	4.6	4.3	5.3	4.3	5.3	5.3

Table 3.12: Answers given to each statement in the questionnaire for each participant.

P	**S1**			**S2**			**S3**			**S4**			**S5**			**S6**			**S7**		
	R	C	CR	R	C	CR	R	C	CR	R	C	CR	R	C	CR	R	C	CR	R	C	CR
1	7	6	5	2	1	2	6	4	3	1	1	2	6	6	5	5	5	5	4	5	4
2	6	5	6	4	4	4	5	5	6	3	1	2	6	6	6	3	6	4	3	6	5
3	6	6	4	6	5	5	3	2	4	3	2	5	5	6	5	2	2	3	2	2	3
4	3	7	7	6	3	4	3	2	3	7	1	1	7	7	7	1	3	1	1	1	1
5	6	5	5	4	4	2	6	5	4	3	4	3	5	5	6	6	3	5	6	3	3
6	4	5	6	6	4	5	4	2	3	6	2	1	5	6	6	2	2	1	3	1	1
7	2	4	4	3	5	6	2	3	3	6	1	4	6	6	6	1	1	1	1	1	1
8	6	6	6	7	5	5	6	5	3	4	5	3	7	7	7	5	3	2	1	1	1
9	5	7	6	6	4	3	5	2	3	6	1	2	6	7	7	4	2	2	5	5	2
10	4	4	5	3	6	4	5	4	6	3	6	6	5	5	4	5	5	5	3	2	4
11	2	5	6	6	6	6	3	3	5	6	4	2	6	6	6	5	3	1	5	2	2
12	7	7	5	2	2	2	5	6	1	1	2	5	7	6	7	1	6	1	2	1	1
13	6	5	6	3	6	1	1	2	2	2	3	2	7	6	6	3	2	2	2	2	2
14	6	6	7	5	2	2	2	2	2	5	5	5	7	7	7	4	4	2	1	1	1
15	3	6	5	5	2	3	5	6	5	5	2	3	5	6	5	5	4	4	5	2	3
16	4	4	5	6	4	6	5	2	4	5	3	4	4	6	5	2	1	3	2	1	1
17	4	5	5	5	3	3	3	2	1	6	3	3	5	6	7	1	1	3	1	1	1
18	4	5	5	6	4	4	4	4	4	5	3	2	5	5	6	5	3	2	5	2	4
19	5	6	3	2	4	4	3	5	5	3	2	5	6	6	6	2	6	5	2	5	5
20	3	3	4	3	6	6	3	2	3	6	7	6	6	1	6	2	7	1	2	4	2
21	6	5	5	5	5	5	2	2	3	4	3	3	6	6	6	2	2	2	2	2	2
22	5	4	4	7	7	7	7	5	6	6	3	3	7	7	7	1	1	1	1	1	1
23	6	7	7	6	7	7	7	6	6	4	4	2	7	7	6	6	3	7	5	5	6
24	4	4	5	5	5	3	4	6	2	6	5	4	5	4	5	6	3	3	6	4	6
25	6	7	5	5	5	3	1	1	3	1	1	2	6	7	7	2	1	2	3	5	7
26	5	2	1	3	6	7	3	1	1	2	7	7	7	3	2	2	1	4	7	1	1
27	6	5	5	3	6	6	5	2	4	2	7	5	6	3	3	5	1	3	2	2	5
28	7	6	7	3	6	2	2	5	1	3	3	2	7	6	7	1	5	3	1	5	5
29	6	4	7	4	6	2	6	6	1	5	6	4	5	5	5	4	6	6	5	6	5
30	3	6	6	5	2	5	5	5	5	5	2	4	6	7	6	2	6	4	2	5	4

S1: I saw a lot of interesting movies
S2: I saw a lot of movies that didn't interest me
S3: I encountered a lot of movies that I didn't expect to see
S4: I thought it was hard to find movies I wanted to watch
S5: I was satisfied with the movies I picked for my fictional flight
S6: The movies I picked for my fictional flight were different from what I expected to pick
S7: The movies I picked for my fictional flight were different from what I normally watch

Table 3.13: Interview questions and answers about the specific arrangements.

Question	Answer	#	Explanation	#
Which arrangement of the movies did you prefer?	R	5	Everything is mixed up, so it doesn't narrow your vision	3
			I first though it'd be very annoying, but it was surprisingly nice	1
			I didn't have to think about the structure, I could just browse	1
			If I'm not looking for a specific movie I prefer to see all sorts of things	1
	C	10	I could easily find movies I wanted to see	4
			I liked navigating by genre/genre labels were very convenient	3
			I found more movies that I wanted to watch than with the other settings	2
			Everything seemed to be in the right place	2
			The movies that I wanted to see were close together	1
			I didn't encounter many uninteresting movies	1
			You can search more directly	1
	CR	2	It felt like I could find what I wanted most easily with this setting	1
	C, CR	10	Didn't see the difference between C and CR	7
			With this setting I can avoid genres/movies that don't interest me	4
			The sorting helped me navigate through the wall	3
			I like to have a starting point/direction	2
			It was easy to switch between genres	1
			It was easier to find a movie I liked and their related movies	1
			I really liked that I already knew the genre of movies to some degree	1
	No preference	3	It depends, with something specific in mind I would use C or CR, with no specific intent I would use R	2
			It depends, I'd use C or CR for quickly finding something, if I had more time I'd use R	1
Which arrangement of the movies did you prefer least?	R	20	It was harder to search for something specific than with the other settings	7
			I saw a lot of movies that didn't interest me	4
			I couldn't easily find movies I wanted to watch	3
			I wasn't in control of the movies I saw	3
			Searching based on genres and that was more convenient with the other settings	2
			I missed a starting point, e.g. a movie or genre that I found interesting	2
			Too chaotic	2
			Too many movies	1
	C	3	It narrows your vision because all the same movies are together	2
			I couldn't really find movies I wanted that I hadn't already found with CR	1
			I felt like I often saw the same movies	1
			I was thinking about the structure	1
	CR	4	Sorting seemed to make less sense/random movies were annoying	2
			It wasn't as easy to find movies with this setting	1
			I didn't have the feeling that I saw a lot of new movies	1
	No preference	3	See explanation above	
Is there a different arrangement of the movies that you think would be better that didn't appear in this test?			Besides genre, also on rating	7
			Add movies to all genre clusters in which they belong	2
			If a movie has multiple genres, put it between the clusters it belongs to	1
			On popularity/rating, with the highest rated/most popular movies in the center	1
			Sort on the director/writer of the movie	1
			On colour, to pinpoint movies of which you know what the poster looks like	1
			Write/country of story origin	1

Table 3.14: Interview questions and answers about the interface in general.

Question	Answer	#
What did you like about the interface?	I saw a lot of movies (in a short time)	10
	It works smoothly	8
	I liked the filters	8
	I encountered a lot of surprising/unexpected/interesting movies	7
	I liked the use of movie posters	6
	It was easy to use/intuitive	5
	I liked that you could zoom	4
	I liked the way of browsing	4
	It looks good	4
	More fun than scrolling through a list	3
	It can handle a lot of movies (without becoming chaotic)	2
	Very efficient use of screen space	2
	Easy to find (new) movies	2
	Enough information for picking a movie	2
	Clear overview	2
	It worked like I expected it to work	2
	I liked the grid structure	2
	It was easy to find related movies	2
	Easier to use than existing movie browsing interfaces	1
	I was able to find a lot of movies I wanted to watch	1
What didn't you like?	You can only judge by poster (at the first sight)	5
	I tended to only pick movies I was already familiar with (because I had to judge on the poster)	2
	Only suited for touch screen	1
	Movie centered on click. This made me lose my orientation sometimes.	1
	If you have a specific movie in mind there is no easy way to find it	1
	I couldn't click on movies excluded by the filter	1
	With my dyslexia is was sometimes hard to to work with small posters	1
	It's a little chaotic/intense	1
	I didn't like being able to scroll both ways, I'd prefer scrolling only one way	1
What would you add/improve upon?	Search bar/function	14
	Keep filters on when switching between movies or closing the panel	9
	Add an easy way to show all movies of a certain collection (e.g. Star Wars)	3
	Movie description more in focus so you don't always need to scroll	1
	Being able to only show cluster labels that you're interested in	1
	More information to the panel	1
	Filters that you can use independent from the selected movies	1
	Double tap to zoom	1
	Sometimes I lost track of where I was. Highlight selected movie more	1
	Highlight everything in the cluster if you tap on a cluster label	1
	Somehow include recommendations	1
	Indicate the number of movies highlighted with the current filter(s)	1
	Filter out certain movies before you start	1

Table 3.15: Interview questions and answers about the interface in general.

Question	Answer	#	Explanation	#
Did you enjoy using the interface?	Yes	24	The interface was innovative/new/surprising	5
			It worked pleasantly	4
			It was a fun experience	3
			I liked seeing a lot of movies	3
			More enjoyable than existing movie browsing interfaces (e.g. Netflix)	3
			Nice just browsing through all the movies	2
			You don't get handed on a platter what you have to watch	2
			I saw a lot of movies I wanted to watch	2
			Filters were fun	2
			You can find all sorts of movies	1
			It gave me inspiration for movies to watch	1
			It was convenient to find something you'd normally wouldn't really watch	1
			Matched my need better than existing movie browsing interfaces	1
			Lot of freedom to browse	1
			Better than browsing through a list	1
	Neutral	4	I wasn't ecstatic to use it or anything, but it worked fine	4
	No	2	Browsing/looking for movies isn't my thing	2
Would you use this interface again?	Yes	26	It works better for me than existing movie browsing interfaces (e.g. Netflix)	5
			Because I never know what to pick	3
			Because it gives a lot of options	2
			Because I found a lot of movies	1
			Because it works pleasantly	1
			When I have no idea what to watch/not looking for anything specific	10
			If I had to pick multiple movies at once, e.g. for a vacation or flight	3
			When I want to see a movie I don't know	3
			If it had all the movies a streaming services had to offer I would use it	2
			If I want to have a lot of control over what movies I see	1
			When picking a movie with roommates	1
			When I'm open to anything	1
			If I want to watch a movie with a certain genre, and I want to see my options	1
	No	4	I don't watch movies on mobile devices	1
			I don't like looking for movies	1
			I generally already know what I want to watch, so I have no need for this	1
			An interface like Netflix' works better for me, I like the recommendations	1

4 Demonstration Interface

After conducting both the pilot study and a second, more elaborate user study, we improved the interface according to our findings and finally present a demo of the interface concept. Note how we still call it a “concept”, as it is above all about the high level concepts, and less about the exact implementation of each feature in the interface. For example, we do not claim that our visualization of the filter functionality is the only and optimal way to do so, instead we plead that having filters which allow users to easily combine metadata to search for related movies is a beneficial feature. This section contains a more elaborate description of the interface presented in the demo paper in Chapter 2.

4.1 Implementation

The demo interface is implemented as an Android application and designed for use on touch-screen tablets. All movie related data comes from The Movie Database (TMDb), using their API. Hence, the screenshots in this section contain movie posters that might be copyright protected. These were downloaded from *themoviedb.org* and used under the fair-usage policy which applies to copyright protected images.

Movie wall The core element of the demo interface (and its previous versions) is the large grid filled with movie posters. It consists of 2160 movies (72 by 30), which are the top movies from TMDb ranked by popularity. This number did not seem to overwhelm users in the second study and is thus kept the same for the demo. However, we are confident number can be changed, as long as the resulting movie collection is interesting and relevant for its user. This *movie wall* can be navigated using standard touch gestures: multi-touch pinch-to-zoom and single touch dragging. The most zoomed out and zoomed in states of the grid are shown in the left and right images in Figure 3.5, respectively. Completely zoomed out the entire grid is visible, meant to give the user an overview of the size and structure of the movie collection. Completely zoomed in, every movie poster should be easily recognisable for the user.

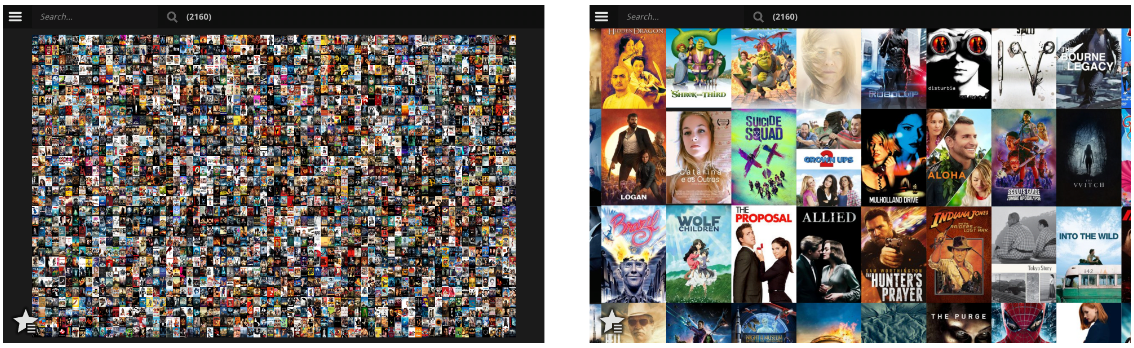

Figure 3.5: The most zoomed out (left) and zoomed in (right) states of the demo interface. In the screenshots shown, the movies are randomly distributed over the wall. All movie posters were downloaded from www.themoviedb.org and used under the fair usage policy.

Initial zoom level The initial zoom level is shown in Figure 3.6. This is the same zoom level as used in the second version of the interface. A completely zoomed out start did not work very well, as it was quite overwhelming for the user. Moreover, it generally caused them to zoom in relatively far, to what we assumed more commonly used image sizes, and thus not even trying smaller zoom levels. The latter is not a direct issue, but it prevents the user from taking full advantage of their capabilities to recognise smaller images and thus the interface design too.

During the second user study, a larger zoom level was used. We saw that users had much less tendency to zoom in very far, and on average spent one-fifth of their time on or close to this zoom level. The results of the second study indicate that this initial zoom level worked quite well, hence it was kept the same for the demo interface.

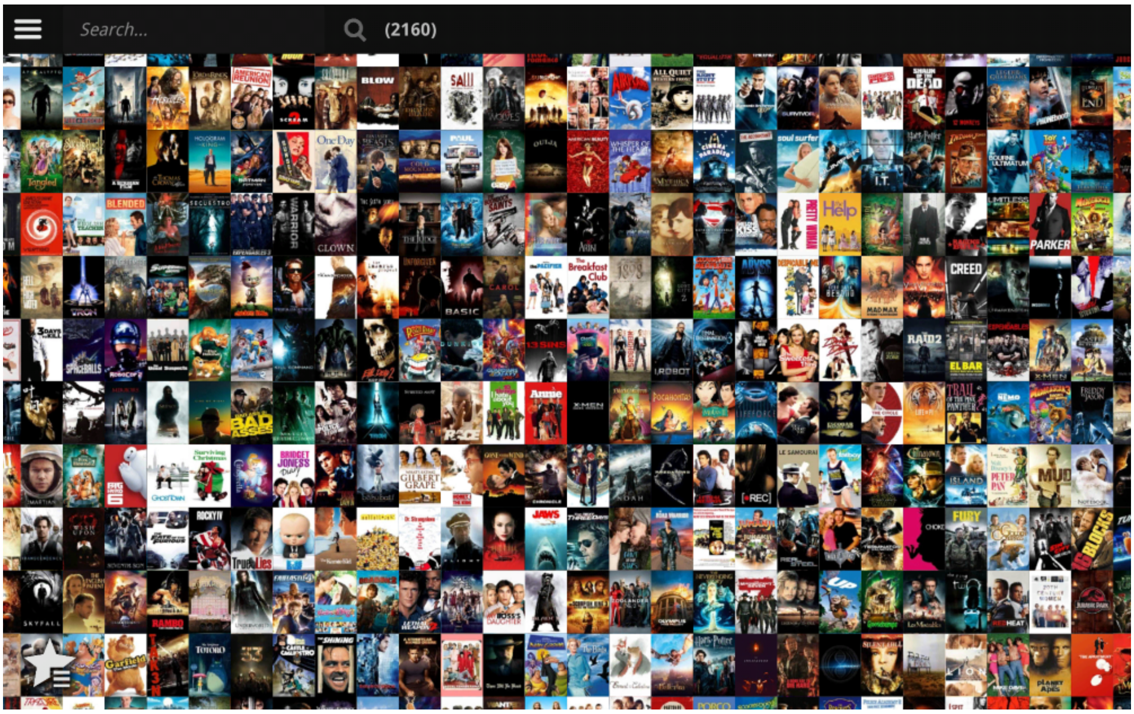

Figure 3.6: The initial zoom level of the demo interface. In the screenshots shown, the movies are randomly distributed over the wall. All movie posters were downloaded from www.themoviedb.org and used under the fair usage policy.

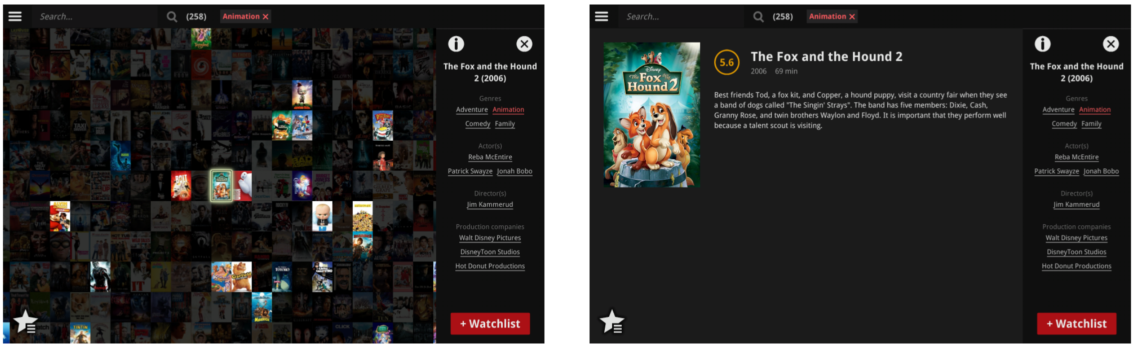

Figure 3.7: Both images show the side panel containing all movie data that can be applied as a filter. The right images shows the panel with room for more detailed material, which can be opened using the information button in the side panel. All movie posters were downloaded from www.themoviedb.org and used under the fair usage policy.

Movie details Selecting a movie highlights it, so the user is less likely to lose its position in the grid, and subsequently opens a panel with movie details, such as genres and actors. In the initial version of the interface, the movie details were put in a side panel which covered only part of the screen space. This allowed the user to still interact with the movie wall with the panel opened. The idea behind this design was that it could open up for more coupling between the wall and the movie details. This coupling was introduced in the second version of the interface with the addition of filters. This caused an increase in interaction with the movie wall with the panel opened, as the filters were a quite successful feature. Still, the panel took up 40 percent of the screen space while opened, which is relatively much. Furthermore, the side panel contained some information that we assumed was not useful to have while interacting with the wall, for example the movie plot. Hence, in the demo interface the movie details are split into two panels, a side panel containing all the filters and the other more detailed information about the movie. The side panel takes up only 20 percent of the screen space. The second panel is shown instead of the movie wall, since we assume it is rarely useful to look at both simultaneously. This design also creates more room for additional material related to a movie, for example a trailer (not yet implemented in the demo). Both panels are shown in Figure 3.7.

Filters The genres, actors, directors and production companies in the side panel can all be applied as filters. Filters highlight everything on the movie wall that meet the currently specified filter criteria. Multiple filters can be combined, allowing users to easily search for movies based on multiple criteria, without having to enter complex queries. The second study showed this was a successful feature, and many participants explicitly expressed their appreciation for the filters.

An issue with the filters in the previous version of the interface was that they would turn off as soon as the user switched movies or closed the details of the current movie. This was undesirable, as it did not allow the user to browse with filtering enabled with the panel closed. Furthermore, selecting all filters again was a tedious process in case a user wanted to keep browsing based on the same criteria after switching movies. To solve this issue, every filter added in the demo interface is not only marked in the side panel, but additionally added to the top of the screen, shown in Figure 3.8. To turn a filter off again, it can either be deselected in the panel, or removed from the top of the screen (when the movie details are closed or changed). Filters have different colours to indicate the different types of filters (genre, cast, director, company).

The demo interface features a search bar as well, which allows the user to specify custom filters. While this will cost users slightly more effort, it grants them the freedom to add criteria not (currently) provided by the standard filters. In the demo interface, a custom filter returns any movie containing the keyword in the following data: (original) title, tagline, characters, cast, crew, genres and companies. An example of a custom filter is shown in Figure 3.8. This was a feature highly requested in the second user study.

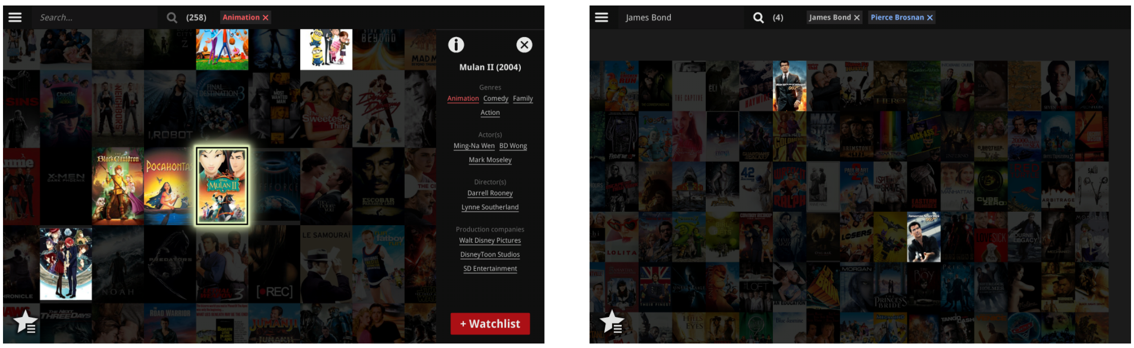

Figure 3.8: The left image shows the movie wall with the single genre filter *Animation* applied. The right images shows multiple filters: *Pierce Brosnan* and *James Bond*, a combination that highlights all movies starring the actor *Pierce Brosnan* and containing the keyword *James Bond*. *James Bond* is a custom filter in this case. All movie posters were downloaded from www.themoviedb.org and used under the fair usage policy.

Arrangement of movies In the second user study we compared different arrangement of movies. However, these results did not show any significant differences in user experience or in satisfaction or number of movies found. The main difference was in user preference: most preferred a clustered arrangement over a random one. Hence, we propose to have at least a default, clustered option for such an interface. The clustered arrangement is shown in Figure 3.9. However, one-sixth of the participants had a strong preference for the random arrangement, which is why we propose to include it. In the demo interface users can choose between either one.

Furthermore, we hypothesize that the disfavour of the random arrangement was perhaps not mainly because the movies were randomly distributed, but rather because of the inability to avoid uninteresting content. Therefore, we believe that an improvement for the random arrangement could be to let users pre-filter the wall, for example excluding movies with specific genres or other metadata. However, this does require extra effort from a user's side.

A semi-clustered arrangement did not work very well in the second study. It did not perform worse either, but the study did not prove any of the assumptions that it would be a good middle ground between clustered and random. Hence, such an arrangement would need more investigation to prove its usefulness. Moreover, some users saw the randomization as a flaw in the system, which is undesirable.

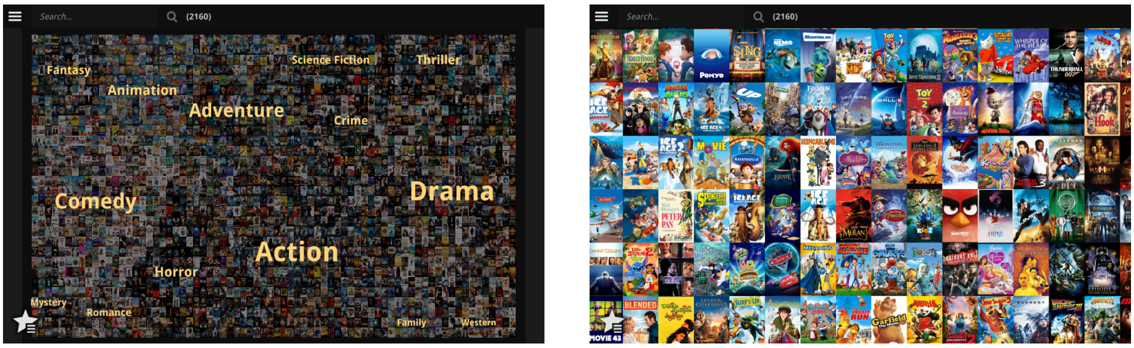

Figure 3.9: Arrangement of movies based on genre. The left image shows it completely zoomed out, with labels that indicate the position and size of each cluster. The right image shows the same arrangement, zoomed in on a part of the animation cluster. All movie posters were downloaded from www.themoviedb.org and used under the fair usage policy.

Conclusions and Future Work

In this thesis project we have investigated and verified the potential, feasibility and usefulness of the interface concept we call the *MovieWall*. It facilitates visual, exploratory browsing through large movie collections. An informal pilot study in the early stages of the project showed the potential of the concept and revealed issues with the design. We improved the interface with these issues in mind. A second, more detailed user study was conducted with the improved interface. It showed that the issues with the first version were largely tackled. Analysis of the interaction data gathered during the usage of the interface also further confirmed the willingness of users to browse with relatively small images. The fact that humans are very capable of recognising small images was already shown in earlier research, but observation that they are also willing to do so, is new. The subjective feedback given by the participants was overwhelmingly positive, 80% of all participant enjoyed using the interface and 87% would use such an interface again. Based on the findings in both users studies, we proposed a demonstration implementation of the interface concept.

For future work, it would be interesting to integrate this design in existing movie browsing interfaces, as it is meant as a complementary, alternative visualization. Our work strongly suggests there is a demand for an interface like the *MovieWall* (or the functionality it provides). Yet, our user studies did not explicitly test whether users would choose this visualization over the standard one (in certain situations), besides subjective statements. So analysing its usage in a more realistic setting, could further prove its usefulness.

Our research indicated there is a desire to search based on various movie criteria, while this is not a feature often encouraged or focussed on in standard browsing interfaces. When provided with filters for genre, cast, crew and production companies, each filter type was used to a certain extent. Future research could look into what sorts of metadata users would use to navigate by, and better ways to incorporate this in browsing interfaces.

To test the potential of the interface, we chose to conduct our user studies with young, early adopters (ages between 20-30), likely quite familiar with both touchscreen devices and movie browsing applications like Netflix. Now we have verified the usefulness of the interface for this group, it would be good to expand these tests to a wider variety of potential users. One example would be middle-aged or even elderly people. Older people generally have more difficulty to focus on objects close by, which is potentially an issue for the usability of such an interface.

We restricted our experiments to touchscreen tablets. For future work, it would be interesting to test its feasibility on other devices as well, for example smartphones or computers. Smartphones generally have rather small screens and it would need additional testing to see whether the *MovieWall* interface is still suitable in this case. The interface might even have potential for televisions, since they often have large screens. We do think the standard television remotes are not suitable to control such an interface. However, nowadays we also see alternatives to control televisions, like game controllers and smartphones, which can potentially support the more complex interaction needed to operate such an interface.

Finally, we are confident this interface concept cannot only be applied to movie browsing, but to other (visual) domains as well, as it can essentially provide a visually appealing start to the exploration of any collection.

References

[1] David Ahlström et al. "A user study on image browsing on touchscreens". In: *Proceedings of the 20th ACM international conference on Multimedia*. ACM. 2012, pp. 925–928.

[2] Kai Uwe Barthel, Nico Hezel, and Radek Mackowiak. "Graph-based browsing for large video collections". In: *International Conference on Multimedia Modeling*. Springer. 2015, pp. 237–242.

[3] James Bennett, Stan Lanning, et al. "The netflix prize". In: *Proceedings of KDD cup and workshop*. Vol. 2007. New York, NY, USA. 2007, p. 35.

[4] Patrick Chiu et al. "MediaMetro: browsing multimedia document collections with a 3D city metaphor". In: *Proceedings of the 13th annual ACM international conference on Multimedia*. ACM. 2005, pp. 213–214.

[5] Patrick Chiu et al. "mTable: browsing photos and videos on a tabletop system". In: *Proceedings of the 16th ACM international conference on Multimedia*. ACM. 2008, pp. 1107–1108.

[6] Ork De Rooij, Cees GM Snoek, and Marcel Worring. "Balancing thread based navigation for targeted video search". In: *Proceedings of the 2008 international conference on Content-based image and video retrieval*. ACM. 2008, pp. 485–494.

[7] Ork De Rooij, Cees GM Snoek, and Marcel Worring. "Mediamill: semantic video search using the rotorbrowser". In: *Proceedings of the 6th ACM international conference on Image and video retrieval*. ACM. 2007, pp. 649–649.

[8] Nuno Gil et al. "Going through the clouds: search overviews and browsing of movies". In: *Proceeding of the 16th International Academic MindTrek Conference*. ACM. 2012, pp. 158–165.

[9] Ai Gomi and Takayuki Itoh. "Mini: A 3d mobile image browser with multi-dimensional datasets". In: *Proceedings of the 27th annual ACM Symposium on Applied Computing*. ACM. 2012, pp. 989–996.

[10] Marco A Hudelist, Klaus Schoeffmann, and David Ahlstrom. "Evaluation of image browsing interfaces for smartphones and tablets". In: *Multimedia (ISM), 2013 IEEE International Symposium on*. IEEE. 2013, pp. 1–8.

[11] Marco A Hudelist, Klaus Schoeffmann, and Laszlo Boeszoermenyi. "Mobile video browsing with a 3D filmstrip". In: *Proceedings of the 3rd ACM conference on International conference on multimedia retrieval*. ACM. 2013, pp. 299–300.

[12] Wolfgang Hürst et al. "Keep moving!: revisiting thumbnails for mobile video retrieval". In: *Proceedings of the 18th ACM international conference on Multimedia*. ACM. 2010, pp. 963–966.

[13] Wolfgang Hürst et al. "Size matters! how thumbnail number, size, and motion influence mobile video retrieval". In: *Advances in Multimedia Modeling* (2011), pp. 230–240.

[14] Yanir Kleiman et al. "Dynamicmaps: similarity-based browsing through a massive set of images". In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM. 2015, pp. 995–1004.

[15] Yanir Kleiman et al. "Toward semantic image similarity from crowdsourced clustering". In: *The Visual Computer* 32.6-8 (2016), pp. 1045–1055.

[16] Yehuda Koren. "The bellkor solution to the netflix grand prize". In: *Netflix prize documentation* 81 (2009), pp. 1–10.

[17] Janette Lehmann et al. "Models of user engagement". In: *International Conference on User Modeling, Adaptation, and Personalization*. Springer. 2012, pp. 164–175.

[18] Hao Liu et al. "Effective browsing of web image search results". In: *Proceedings of the 6th ACM SIGMM international workshop on Multimedia information retrieval*. ACM. 2004, pp. 84–90.

[19] Thomas Low et al. "Exploring Large Movie Collections: Comparing Visual Berrypicking and Traditional Browsing". In: *International Conference on Multimedia Modeling*. Springer. 2017, pp. 198–208.

[20] Thomas Low et al. "Visual berrypicking in large image collections". In: *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*. ACM. 2014, pp. 1043–1046.

[21] Pedro Martins, Thibault Langlois, and Teresa Chambel. "MovieClouds: content-based overviews and exploratory browsing of movies". In: *Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments*. ACM. 2011, pp. 133–140.

[22] Christopher Müller, Martin Smole, and Klaus Schöffmann. "A demonstration of a hierarchical multi-layout 3D video browser". In: *Multimedia and Expo Workshops (ICMEW), 2012 IEEE International Conference on*. IEEE. 2012, pp. 665–665.

[23] Heather L O'Brien, Rafa Absar, and Helen Halbert. "Toward a Model of Mobile User Engagement". PhD thesis. University of British Columbia, 2013.

[24] Heather L O'Brien and Elaine G Toms. "Examining the generalizability of the User Engagement Scale (UES) in exploratory search". In: *Information Processing & Management* 49.5 (2013), pp. 1092–1107.

[25] Heather L O'Brien and Elaine G Toms. "The development and evaluation of a survey to measure user engagement". In: *Journal of the American Society for Information Science and Technology* 61.1 (2010), pp. 50–69.

[26] Zoran Pečenovió et al. "Integrated browsing and searching of large image collections". In: *International Conference on Advances in Visual Information Systems*. Springer. 2000, pp. 279–289.

[27] Kerry Rodden et al. "Does organisation by similarity assist image browsing?" In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM. 2001, pp. 190–197.

[28] Kerry Rodden et al. "Evaluating a visualisation of image similarity as a tool for image browsing". In: *Information Visualization, 1999.(Info Vis' 99) Proceedings. 1999 IEEE Symposium on*. IEEE. 1999, pp. 36–43.

[29] Gerald Schaefer. "A next generation browsing environment for large image repositories". In: *Multimedia Tools and Applications* 47.1 (2010), pp. 105–120.

[30] Gerald Schaefer. "Interacting with image collections: visualisation and browsing of image repositories". In: *Proceedings of the 20th ACM international conference on Multimedia*. ACM. 2012, pp. 1527–1528.

[31] Klaus Schoeffmann and David Ahlström. "An evaluation of color sorting for image browsing". In: *International Journal of Multimedia Data Engineering and Management (IJMDEM)* 3.1 (2012), pp. 49–62.

[32] Klaus Schoeffmann and David Ahlström. "Using a 3d cylindrical interface for image browsing to improve visual search performance". In: *Image Analysis for Multimedia Interactive Services (WIAMIS), 2012 13th International Workshop on*. IEEE. 2012, pp. 1–4.

[33] Klaus Schoeffmann, David Ahlstrom, and Christian Beecks. "3d image browsing on mobile devices". In: *Multimedia (ISM), 2011 IEEE International Symposium on*. IEEE. 2011, pp. 335–336.

[34] Klaus Schoeffmann, David Ahlström, and Laszlo Böszörmenyi. "3d storyboards for interactive visual search". In: *Multimedia and Expo (ICME), 2012 IEEE International Conference on*. IEEE. 2012, pp. 848–853.

[35] Klaus Schoeffmann, David Ahlström, and Laszlo Böszörmenyi. "A user study of visual search performance with interactive 2d and 3d storyboards". In: *International Workshop on Adaptive Multimedia Retrieval*. Springer. 2011, pp. 18–32.

[36] Klaus Schoeffmann, David Ahlström, and Marco A Hudelist. "3-D interfaces to improve the performance of visual known-item search". In: *IEEE Transactions on Multimedia* 16.7 (2014), pp. 1942–1951.

[37] Klaus Schoeffmann and Laszlo Boeszoermenyi. "Image and video browsing with a cylindrical 3D storyboard". In: *Proceedings of the 1st ACM International Conference on Multimedia Retrieval*. ACM. 2011, p. 63.

[38] Klaus Schoeffmann and Manfred del Fabro. "Hierarchical video browsing with a 3D carousel". In: *Proceedings of the 19th ACM international conference on Multimedia*. ACM. 2011, pp. 827–828.

[39] Grant Strong, Orland Hoeber, and Minglun Gong. "Visual image browsing and exploration (Vibe): User evaluations of image search tasks". In: *International Conference on Active Media Technology*. Springer. 2010, pp. 424–435.

[40] Damien Tardieu et al. "Browsing a dance video collection: dance analysis and interface design". In: *Journal on Multimodal User Interfaces* 4.1 (2010), pp. 37–46.

[41] Antonio Torralba, Rob Fergus, and William T Freeman. "80 million tiny images: A large data set for nonparametric object and scene recognition". In: *IEEE transactions on pattern analysis and machine intelligence* 30.11 (2008), pp. 1958–1970.

[42] Andreas Töscher, Michael Jahrer, and Robert M Bell. "The bigchaos solution to the netflix grand prize". In: *Netflix prize documentation* (2009), pp. 1–52.

[43] Ryen W White and Resa A Roth. "Exploratory search: Beyond the query-response paradigm". In: *Synthesis Lectures on Information Concepts, Retrieval, and Services* 1.1 (2009), pp. 1–98.