Exploring Visualization Challenges for Interactive Recommender Systems

MANDY KECK, DIETRICH KAMMER, Technische Universität Dresden

Users are faced with an increasing information overload problem in large, complex data collections. Recommender systems reduce the data set to a manageable size by providing suggestions to the user. Research in the last years has primarily focused on the quality of the underlying algorithms. Recent research started to focus on the user experience in recommender systems. The main challenges are transparency, controllability, explorability, and context-awareness. Interactive visualizations have the potential to address all of these issues. In this paper, we present three user interface concepts for different usage scenarios: movie, activities, and travel search. We propose a taxonomy of user interface building blocks to evaluate these concepts with regards to the visualization challenges.

CCS Concepts: • **Information systems → Recommender systems**; Search interfaces; • **Human-centered computing** → *Interaction paradigms*; *Information visualization*;

Additional Key Words and Phrases: Recommender Systems, Information Visualization, Human Computer Interaction

1 INTRODUCTION

In large, complex data collections such as product catalogues, users are faced with an increasing information overload problem and it is difficult to find suitable items. To this end, recommender systems offer an appropriate mechanism to reduce the data set to a manageable size by identifying items that may be interesting for the particular user [15]. During the last years, the quality of these systems has been considerably improved. However, most research has primarily focused on the algorithms to enhance performance and accuracy [13]. Latest research started to go beyond the optimization of algorithms by focusing on the user experience with recommender systems. Previous studies have shown that visual features and enhanced interaction improve the user engagement with the system and the acceptance of the recommended items [14]. Our literature review revealed many issues that visualizations can help to solve.In this paper, we focus on the following issues with regards to recommender systems that are easy to use for end users and support them in exploring and understanding the information space:

- **Transparency**. Supporting the user in understanding the reasons behind the recommendations
- **Controllability**. Providing user control over features that influence the recommender algorithm
- **Explorability**. Presenting visualizations to browse the entire information space, e.g. related items that are not recommended
- **Context-Awareness**. Considering different situations, such as mood, time, individual, or collaborative scenarios

In this paper, we present three different user interface concepts tackling these issues by interactive visualizations. The paper is structured as follows: the next section considers related work regarding to the introduced visualization challenges. Subsequently, we illustrate interface concepts for recommender systems in three different scenarios:

Author's address: Mandy Keck, Dietrich Kammer, Technische Universität Dresden, Dresden, Germany, 01187, {firstname.lastname}@tu-dresden.de.

VisBIA 2018 – Workshop on Visual Interfaces for Big Data Environments in Industrial Applications. Co-located with AVI 2018 – International Conference on Advanced Visual Interfaces, Resort Riva del Sole, Castiglione della Pescaia, Grosseto (Italy), 29 May 2018

movies, activities, and travel search. Each concept will highlight visualization and interaction techniques to address the challenges. In the last sections, we will discuss these approaches and give an outlook of future work.

2 RELATED WORK

User experience in recommender systems can be considerably improved using interactive visualizations [14]. Visualization challenges identified in the literature range from transparency to diversity (cp. [18], [9]). In this paper, we address the main challenges transparency, controllability, explorability, and context-awareness to increase the user experience of the end user. Challenges such as understanding the inner logic of the recommender system or manipulating the level of diversity are not considered in this paper, because they focus user groups with a higher technical expertise.

Recommender systems lack **transparency**, when they appear as "black boxes" to the user, making it incomprehensible how recommendations are generated and why a specific list of items is presented [15]. One method to improve the transparency (also called *"Justification"* in [9]) of a recommender system and the users' trust in the results, are explanations. They can help users to understand the reason behind a recommendation, increase the user's sense of involvement in the recommendation process and can lead to a greater acceptance of the recommender system as a decision aide [11]. One example is *TV Land*, a map-based visualization that depicts the relationships between the most watched TV shows [8]. It highlights TV shows the user watched already and explains the context for new recommendations. Other approaches visualize the user model in order to improve the user's understanding of how his preferences are represented within the system [12].

Further research focuses on the possibilities to influence the recommendation algorithms and increase the **controllability** over the recommendation process by providing the user control over the user model and the personalization [3], by adjusting the influence of different resources [6, 14] or by giving immediate feedback about the provided items to influence the recommender algorithm. The latter is shown in an example that uses a 3D map-based visualization, where users can give feedback in a playful manner and thus manipulate their underlying profile: They can shape the landscape by creating hills and valleys to express their regions of interest and to influence the recommender algorithm [13].

A third aspect affects the **explorability**. The widely used presentation of results in the form of ranked lists impedes getting an overview of the item space and hide relations to unrecommended items. Whilst in the field of information retrieval and information visualization various methods exist for visualizing large dataset such as document collections [10, 16, 17], effective support through visualization in recommender system is still a neglected field of research [13]. The introduced example *TV Land* also supports explorability by providing all TV shows on a map, sorted according to their similarity. Recommended shows are highlighted but still allow the exploration of similar results.

Furthermore, it is important to incorporate contextual information into the recommendation process in order to recommend items to users in certain circumstances [2]. Adomavicius et al. argue that it is important to take the aspect **context-awareness** into account when providing recommendations and distinguish between four context types that should be considered in recommender systems: Physical context (e.g. time, position, activity of the user), Social context (e.g. alone or in a group), Interaction media context (e.g. used device, type of media), and Modal context (e.g. mood, experience and goals of the user) [1]. Bogdanov et al. allow users to configure avatars to express their preferences. To support different user situations while listening to music, they suggest context-dependent avatars, which can be used for listening to recommendations depending on the context (e.g. in the car, at a party) [4].

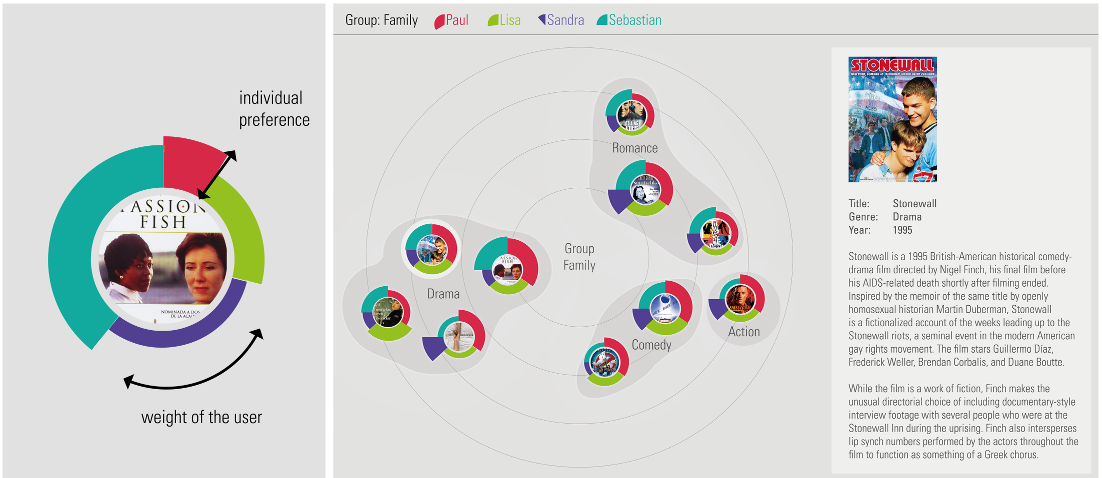

Fig. 1. Structure of a glyph (left), interface with distribution of the recommended movies visualized as glyphs on the left side, details on demand on the right side (right)

3 MOVIE SCENARIO

Most recommender systems focus on recommendations for individual users, neglecting their social context. In many scenarios such as choosing a movie from a movie database, multiple users are involved in the decision process. In this example, we focus on a collaborative search scenario at home, involving a group of friends or a family that is choosing a movie.

3.1 General Concept

The application starts with a configuration view that allows adding new groups and editing existing groups by adding or removing users. As soon as all group members are chosen, each person is represented by his name and a unique color. Furthermore, different priorities can be assigned to selected group members. This might be necessary, when a family with children involved wants to choose a movie. Then the age and preferences of the children have a higher influence on the recommendation algorithm, excluding movies with certain age restrictions.

After the group is selected and adjusted, the recommendations are calculated based on the WARP (Weighted Approximate-Rank Pairwise) algorithm [19]. The left side shows the recommendations using an interactive visualization. The right side offers details on demand after selecting a movie on the left side (see Figure 1, right). The presented features have been implemented in a prototype[1] that is based on the movielens data set[2]. The prototype is implemented in Python and Javascript, using the libraries D3[3] and Flask[4].

[1]video of the prototype: https://www.youtube.com/watch?v=FjlBdS2Gd8I, retrieved on 04.05.2018
[2]https://movielens.org, retrieved on 04.05.2018
[3]https://d3js.org, retrieved on 04.05.2018
[4]http://flask.pocoo.org, retrieved on 04.05.2018

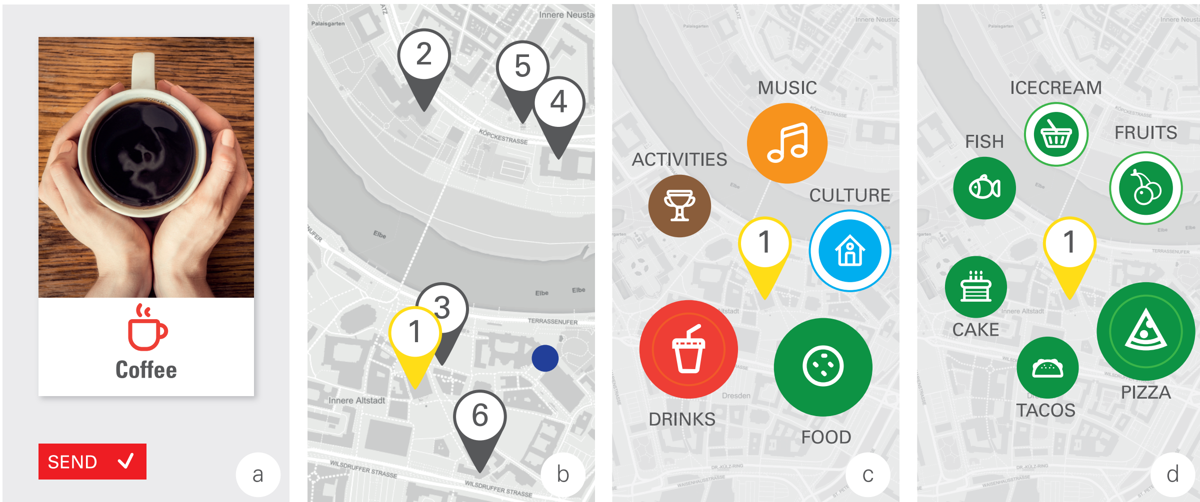

Fig. 2. Preference dialog (a), Map visualization (b), Location selection (c), and subcategory configuration (d)

3.2 Visualization Challenges

The movies are arranged in a radial layout. Movies with the best average rating of all group members are placed close to the center. Hence, the **social context** is visualized. Additionally, they are clustered according to their genre tags, which is represented by a transparent layer surrounding the cluster. In the data set, different genres are assigned to each movie. The genre with the highest value specifies in which cluster the movie is located in the initial view (see Figure 1, right). **Explorability** is made possible by offering zoomable clusters. If the group members decide for movies with the genre tag "comedy", they can select the particular cluster to zoom in. After that, more movies for this cluster are shown in a similar view, but filtered by the tag "comedy". Each movie is represented as a glyph [5], with a preview of the movie cover in the center (see Figure 1, left). Each segment of the ring glyph represents one user in the corresponding color. The width of each segment shows the weight of the user, whereas the height represents how much the particular movie fits to the individual user preferences. The ring glyph provides insight on the impact of each user on the recommendation algorithms. The concept affords **transparency** of the recommendations by visualizing an average rating (position) as well as an individual rating for each user (glyphs). Furthermore, **controllability** is provided, by adding different users to the group and adjust their priorities within the group.

4 ACTIVITIES SCENARIO

The use case in this scenario is the search for activities while walking through a potentially unknown city. Mobile apps are perfect to provide recommendations for locations, events, or restaurants on the go. Every step of the recommendation process is supported by suitable visualizations that are explained in the next subsections.

4.1 General Concept

Users first provide explicit feedback about their preferences using a dialog based card game comparable to the popular dating app "Tinder" (see Figure 2a). The following categories are provided in our prototype of the application, which can also be extended using open linked data: Food (e.g. Pizza, Pasta, Sushi), Drinks (e.g. Beer, Cocktails, Wine), Music (e.g. Electronic, Hip-Hop, Jazz, Rock), Culture (e.g. Museums, Sights, Theatre, Cinema), Activities (e.g. Dancing, Billiard, Dart).

Users interact with the main category cards by swiping left, right, or up. A swipe to the left is a *dislike*, which expresses little interest in a category. By swiping right, this category is added to the user interests as a *like*. Swiping up expresses a *superlike*, which means that this category is crucial to the user. Consequently, sub-categories are selected after performing a superlike in order to specify this interest in more detail. The dialog can be completed at any time using the "send" button. Otherwise, all main categories are presented to the user randomly. In the current prototype, the dialog must always be completed for each session. However, it is conceivable that the completed preference dialog can be saved and loaded each time a new session is started, e.g. in another city.

The current location of the user is automatically leveraged for the recommendations that are calculated based on the user preferences obtained in the initial dialog. Each location that is retrieved via external APIs such as Google Places, Yelp, or Facebook is ranked according to their related categories. A dislike yields a value of `-1`, no rating `0`, like `1` and superlike `2`. Since a location can have several categories, all values are summarized. Thus, locations that meet more user criteria receive a higher ranking. The prototype[5] is implemented using Node.js[6].

4.2 Visualization Challenges

Recommendations are then visualized on a map as markers with numerical values for their ranking (see Figure 2b). User location is visualized as well in order to visualize the **physical context** (blue dot on the map). The top recommendation is visualized as a yellow marker. Common map interaction such as zooming and panning is supported.

By tapping on a marker, another view for the **transparency** and **controllability** of the recommendation results is displayed (see Figure 2c). In order to make the recommendation more transparent, a combination of a bubble chart and mind-map is used. Simultaneously, the map in the background is frozen and displayed with lower saturation to make the information visualization on top more prominent. Categories are displayed as bubbles in a radial grid around the selected location. Both the user preference and the location properties must be shown in this visualization. Hence, each bubble features two different radii. One radius shows the relevance to the user according to the initial preference dialog. The bubble is displayed with a full color inside this radius. The radius for the location is only displayed by an unfilled circle. Consequently, if the user preference is equal to the location value for this category, the bubble is displayed in full color. If the user preference was higher than the location relevance, the ring representing the location value is visible inside the bubble. If the location has a higher score, a ring is displayed around the user bubble. Tapping a category reveals the subcategories, which are displayed in the same way (see Figure 2d). On larger mobile devices or on desktop computers, it is also possible to display categories and subcategories simultaneously. On smaller screens, a semantic zoom can be used.

Controlling the recommendations is achieved by altering the weight of the selected categories. A tap-and-hold gesture on a node results in a continuous change of its size from small to large. Completing the tap-and-hold gesture results in a new calculation of the recommendations. **Explorability** is provided by panning the map to explore regions further away from the current location of the user or zooming into the map to show more recommendations in this specific area. Through changing the viewport, new recommendations are calculated for this specific area, hence **controllabiliy** is provided as well.

5 TRAVEL SCENARIO

Most travel portals offer the best deals for the desired holiday destination. But if the user does not know exactly, where he or she can go and how to formulate a concrete search query with this vague information need, only

[5]video of the prototype: https://www.youtube.com/watch?v=O7is2SeSeuw, retrieved on 04.05.2018
[6]https://nodejs.org, retrieved on 04.05.2018

Fig. 3. Three parts of the interface: suggested visual concepts (left), sieve visualization technique (middle), and mood board with collected results and trash bin (right)

little guidance is provided. This approach focuses on travel search with vague information need and serves as inspiration and comparison tool for different holiday destinations. The search approach allows a step-by-step reduction of the result set by selecting visualized concepts such as "beach", "culture", and "relaxing". The introduced visualization concept extends the interface "GetInspired" that uses a selection-based recommendation-driven search, based on the principle of divide and conquer [7]. The travel scenario does not apply the recommendation algorithm on the result set directly as shown in the previous examples. In contrast, it suggests concepts that can be used to filter the result set. These recommendations are based on the user's interactions with the system and can be seen as "meta-recommendations" for the final holiday destinations. Hence, this approach is suitable for large and complex taxonomies, that only presents information needed by the user to solve his task.

5.1 General Concept

The interface is divided into three parts. The left-hand side provides nine different visual concepts that describe a holiday destination, such as "warm", "party" and "water sport" (see Figure 3, left part). Based on the selection of the user, nine new concepts are suggested that should lead to a desired result with the fewest number of refinement steps possible. The selection of a visual concept reduces the result set, which is visualized in the center of the interface (see Figure 3, middle). The interface element uses a sieve metaphor that is explained in Figure 4. Each visual concept or filter is represented by a horizontal line, sorted by their selection. The results are represented as vertical lines, which can be extended downwards. In the case that a filter criteria is not matching the particular result, the vertical line is stopped in this position. Hence, just the results that cross all horizontal lines, match all selected filters, which is comparable with a sieve. For this visualization concept, a radial layout is chosen, to use the given space efficiently and to avoid different interpretations that might be triggered through the position of the results. In the radial version (see Figure 3, middle), the filters are sorted from outside to the center of the circle. The results are represented as clockwise ordered lines. Through mouse-over interaction, the user can see previews of the holiday destination in the inner circle of the visualization. Interesting destinations can be collected in the right-hand side of the interface (see Figure 3, right part). This part uses a mood board metaphor and supports the organization and comparison of the collected results. A mood board is a common tool to collect pictures, materials and texts to give a general idea of a topic. Hence, the results are represented as circles. The outer circle serves as menu and the inner circle as preview to present different information about the holiday destination, such as prices and temperatures in the form of a graph and pictures. The elements can be

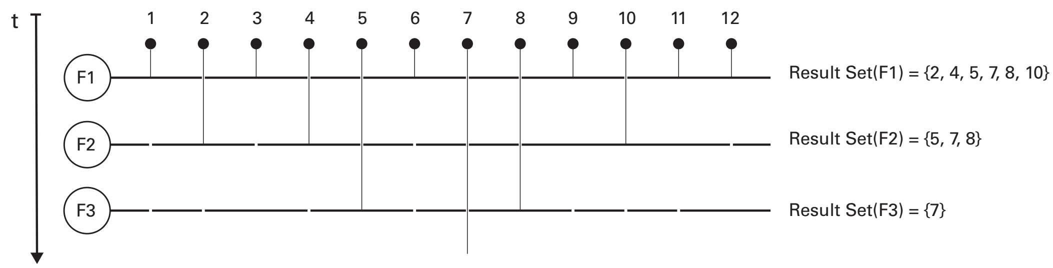

Fig. 4. Sieve metaphor

grouped by placing them close to each other, visualized by a surrounding border that connects these elements and also synchronizes the presented properties of each items. This allows the comparison of the holiday destinations, e.g. finding the cheapest time to travel according to the prices represented in the graph for each month. The prototype[7] is implemented using jQuery[8], D3js and Sylvester[9].

5.2 Visualization Challenges

The sieve metaphor affords **explorability**, since filtered results are not hidden from the user. Rather, he or she is able to hover over the results in order to get explanations about which filter is responsible for the exclusion from the current result set. Hence, **transparency** is supported with respect to the result sets, but not the filters that are influenced by the recommendation algorithm. By removing filters that exclude interesting results, the user is able to reformulate the search query. Hence the recommendation algorithm provides a new set of visual concepts to filter the result set. The inclusion and exclusion of visual elements gives **controllability** over the recommendation algorithm that is recalculated depending on the users' interaction. There is also the possibility to move results to a trash bin in order to influence the recommended filters and results.

6 DISCUSSION

In this paper, we presented a taxonomy of visualization challenges for recommender systems. The presented interface concepts tackle these challenges by using specific building blocks. Figure 5 summarizes the interface building blocks of each usage scenario with regards to **transparency**, **controllability**, **explorability**, and specific aspects of **context-awareness**. In addition, the suitability of each approach with regard to different physical, social, media and modal context are assessed.

The taxonomy is suitable to reflect different solutions with regard to their support for user experience issues in recommender systems. Furthermore, an analysis of more visualization approaches with the help of this taxonomy can provide building blocks to tackle the visualization challenges that serve as an inspiration for future developments.

In terms of **transparency**, glyphs have the potential to show the impact of single preferences on the recommendation algorithms and allow the comparison of the user preferences with the current result set, shown in all three scenarios. Furthermore, the visualization techniques *Radial View* and *Bubble Chart* can support in analyzing and comparing the position of items in the ranking algorithm, whereas the visualization technique *Sieve* supports the user in understanding which preferences match or don't match the viewed result.

[7] video of the prototype: https://www.youtube.com/watch?v=3LixVxYHGPI, retrieved on 04.05.2018
[8] https://jquery.com, retrieved on 04.05.2018
[9] http://sylvester.jcoglan.com, retrieved on 04.05.2018

Challenges	Movie	Activities	Travel
Transparency	Glyphs + Radial View	Glyphs + Bubble Chart	Glyphs + Sieve
Controllability	Group Selection	Tap-&-Hold Gesture	Collect/Exclude Filter Concepts and Results
Explorability	Zoom into Cluster	Pan + Zoom Map	Hightlight Lines
Context-Awareness	Color Coded Users	Map Marker	-
Physical	-	Location	-
Social	Group	Single User	Single User
Media	Tablet + Desktop	Smartphone	Desktop
Modal	vague information need	concrete information need	vage information need

Fig. 5. Taxonomy and evaluation of concepts

Controllability can be applied by adapting the user preferences, illustrated by the tap-and-hold gesture in the activity scenario and the selection and excluding of results in the travel scenario. Also, the feedback on the presented result can offer controllability by collecting interesting or exclude uninteresting results that can influence the recommender algorithm as well, what is shown in the mood board of the travel scenario.

Explorability was also offered in all scenarios. The movie scenario offers a navigation concept that allows the exploration of different regions of interest. The activities scenario supports the well-known interaction techniques panning and zooming to explore different regions of the map. The travel scenario supports a visualization concept that highlights the suggested results but is still showing the filtered items in the context for further exploration and analysis.

The fourth row in the table is related to **context-awareness** and presents interface building blocks that show context-specific information. In the movie scenario, different users are represented by a color-coded glyph segment, whereas the activities scenario presents the user position with a map marker.

The lower part of the table addresses the aspects from Adomavicius et al.[1]. Contextual information can be obtained in different ways, e.g explicitly from the user or implicitly from the environment [9]. Physical context is for example obtained implicitly in the activities scenario, whereas the social context is specified explicitly through the group selection. Media context is provided if the application adapts the appearance of the user interface according to the detected device. This is not implemented in the current prototypes but could be a feature for future work.

Our three demo applications address the modal context with regards to the goals of the user only in a static way. So far, no adaptations of the interfaces are offered if an information need changes from vague to concrete. Both movie and travel scenarios support a vague information need that is mostly associated with offering extensive explorability support in the application. In contrast, the activities scenario is focused on concrete information, which is addressed by the controllability features of the application and the initial preference dialog.

7 FUTURE WORK

From the experience in designing and implementing our demo applications, all three main challenges can be addressed in a given application without sacrificing either transparency, controllability, or explorability. However, there is a rather obvious need for more screen real estate when trying to support explorability or transparency in an interface. Hence, small devices in mobile scenarios usually provide weaker support in these aspects.

The proposed taxonomy is a work-in-progress and should be extended with more building blocks in the future and applied to more solutions. So far, our taxonomy can be seen as a subset of the taxonomy by [9] but is more reduced and focused on the end user with little technical expertise. To prove this assumption, the taxonomy and systems need to be evaluated more thoroughly in future work.

8 ACKNOWLEDGEMENTS

The usage scenarios in this contribution have been developed in a student workshop (Movie scenario: Kristian Kyas and Elisabeth Baudisch, Activities scenario: Julian Haluska and Finn Schlenk) and a diploma thesis (Travel scenario: Marcus Kirsch).

REFERENCES

[1] Gediminas Adomavicius, Bamshad Mobasher, Francesco Ricci, and Alexander Tuzhilin. 2011. Context-Aware Recommender Systems. *AI Magazine* 32, 3 (2011), 67–80. http://www.aaai.org/ojs/index.php/aimagazine/article/view/2364

[2] G. Adomavicius and A. Tuzhilin. 2011. *Context-Aware Recommender Systems*. Springer, Boston, MA.

[3] Fedor Bakalov, Marie-Jean Meurs, Birgitta König-Ries, Bahar Sateli, René Witte, Greg Butler, and Adrian Tsang. 2013. An Approach to Controlling User Models and Personalization Effects in Recommender Systems. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces (IUI '13)*. ACM, New York, NY, USA, 49–56. https://doi.org/10.1145/2449396.2449405

[4] Dmitry Bogdanov, MartÃŋn Haro, Ferdinand Fuhrmann, Anna Xambo, Emilia Gomez, and Perfecto Herrera. 2013. Semantic audio content-based music recommendation and visualization based on user preference examples. *Information Processing & Management* 49, 1 (2013), 13 – 33. https://doi.org/10.1016/j.ipm.2012.06.004

[5] R. Borgo, J. Kehrer, D.H.S Chung, E. Maguire, R.S. Laramee, H. Hauser, M. Ward, and M. Chen. 2013. Glyph-based Visualization: Foundations, Design Guidelines, Techniques and Applications. *Eurographics State of the Art Reports* (May 2013), 39–63. https://www.cg.tuwien.ac.at/research/publications/2013/borgo-2013-gly/

[6] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. TasteWeights: A Visual Interactive Hybrid Recommender System. In *Proceedings of the Sixth ACM Conference on Recommender Systems (RecSys '12)*. ACM, New York, NY, USA, 35–42. https://doi.org/10.1145/2365952.2365964

[7] Andreas Both, Viet Nguyen, Mandy Keck, Dietrich Kammer, Rainer Groh, and Dana Henkens. 2014. Motive-based Search using a recommendation-driven visual divide and conquer approach. *IADIS International Journal on WWW/Internet* 12 (11 2014), 115 – 130.

[8] Emden Gansner, Yifan Hu, Stephen Kobourov, and Chris Volinsky. 2009. Putting Recommendations on the Map: Visualizing Clusters and Relations. In *Proceedings of the Third ACM Conference on Recommender Systems (RecSys '09)*. ACM, New York, NY, USA, 345–348. https://doi.org/10.1145/1639714.1639784

[9] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications* 56 (2016), 9 – 27. https://doi.org/10.1016/j.eswa.2016.02.013

[10] Marti A. Hearst. 2009. *Search User Interfaces* (1st ed.). Cambridge University Press, New York, NY, USA.

[11] Jonathan L. Herlocker, Joseph A. Konstan, and John Riedl. 2000. Explaining Collaborative Filtering Recommendations. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (CSCW '00)*. ACM, New York, NY, USA, 241–250. https://doi.org/10.1145/358916.358995

[12] Yucheng Jin, Karsten Seipp, Erik Duval, and Katrien Verbert. 2016. Go With the Flow: Effects of Transparency and User Control on Targeted Advertising Using Flow Charts. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '16)*. ACM, New York, NY, USA, 68–75. https://doi.org/10.1145/2909132.2909269

[13] Johannes Kunkel, Benedikt Loepp, and Jürgen Ziegler. 2017. A 3D Item Space Visualization for Presenting and Manipulating User Preferences in Collaborative Filtering. In *Proceedings of the 22Nd International Conference on Intelligent User Interfaces (IUI '17)*. ACM, New York, NY, USA, 3–15. https://doi.org/10.1145/3025171.3025189

[14] Denis Parra, Peter Brusilovsky, and Christoph Trattner. 2014. See What You Want to See: Visual User-driven Approach for Hybrid Recommendation. In *Proceedings of the 19th International Conference on Intelligent User Interfaces (IUI '14)*. ACM, New York, NY, USA, 235–240. https://doi.org/10.1145/2557500.2557542

[15] Christian Richthammer, Johannes Sänger, and Günther Pernul. 2017. Interactive Visualization of Recommender Systems Data. In *4th Workshop on Security in Highly Connected IT Systems (SHCIS)*. https://epub.uni-regensburg.de/36211/

[16] Tony Russell-Rose and Tyler Tate. 2013. *Designing the Search Experience: The Information Architecture of Discovery* (1st ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[17] Ben Shneiderman. 2008. Extreme Visualization: Squeezing a Billion Records into a Million Pixels. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data (SIGMOD '08)*. ACM, New York, NY, USA, 3–12. https://doi.org/10.1145/1376616.1376618

[18] Katrien Verbert, Denis Parra, Peter Brusilovsky, and Erik Duval. 2013. Visualizing Recommendations to Support Exploration, Transparency and Controllability. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces (IUI '13)*. ACM, New York, NY, USA, 351–362. https://doi.org/10.1145/2449396.2449442

[19] Jason Weston, Samy Bengio, and Nicolas Usunier. 2011. WSABIE: Scaling Up to Large Vocabulary Image Annotation. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume Three (IJCAI'11)*. AAAI Press, 2764–2770. https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-460