

Interactive recommender systems: a survey of the state of the art and future research challenges and opportunities

Chen He^{a,*}, Denis Parra^b, Katrien Verbert^a

^aDepartement Computerwetenschappen, KU Leuven, Leuven, Belgium

^bPontificia Universidad Católica de Chile Santiago, Chile

Abstract

Recommender systems have been researched extensively over the past decades. Whereas several algorithms have been developed and deployed in various application domains, recent research efforts are increasingly oriented towards the user experience of recommender systems. This research goes beyond accuracy of recommendation algorithms and focuses on various human factors that affect acceptance of recommendations, such as user satisfaction, trust, transparency and sense of control. In this paper, we present an interactive visualization framework that combines recommendation with visualization techniques to support human-recommender interaction. Then, we analyze existing interactive recommender systems along the dimensions of our framework, including our work. Based on our survey results, we present future research challenges and opportunities.

Keywords:

Recommender systems, visualization, user control

1. Introduction

No matter whether we notice it or not, we encounter recommender systems almost everyday, such as Ad recommendations in any possible corner of a web page or product recommendations in online shops. Due to their ability to solve the increasingly severe problem of information overload, recommender systems have gained massive attention over the past decades. The Netflix Prize (Bennett and Lanning, 2007) between 2006 and 2009 and several other challenges organized at the Recommender Systems (RecSys) conference attracted numerous researchers from machine learning and data mining research fields.

Whereas several algorithms have been developed and deployed to suggest relevant items to a user (Adomavicius and Tuzhilin, 2005), there are still several challenges that need to be resolved before recommender systems can realize their full potential. Several recommendation algorithms suffer from *cold start* issues, i.e. they cannot make effective recommendations for new users or for new items that have no explicit or implicit relevance indicators yet (Burke, 2010). In addition, recommender systems often appear as a “*black box*”, i.e. they do not offer the user any insight into the system logic or justification for the recommendations (Sinha and Swearingen, 2002). This black box nature of recommender systems prevents users from comprehending recommended results and can lead to *trust*

issues when recommendations fail (Herlocker et al., 2000). In addition, the approach does not enable users to provide feedback. As the predication of the current interest of the user is often a challenging task, there is a need to develop mixed-initiative approaches that enable users to help steer this process. Such mixed-initiative approaches are also promising to address other issues of recommender systems, such as increasing *diversity* (Hu and Pu, 2011) and *novelty* (Herlocker et al., 2004) of recommended results, and their deployment in *high-risk application domains* such as health-care and financing (McSherry, 2005).

In recent years, researchers have become more aware of the fact that effectiveness of recommender systems goes beyond recommendation accuracy (Swearingen and Sinha, 2001). Thus, research on these human factors has gained increased interest, for instance by combining interactive visualization techniques with recommendation techniques to support transparency and controllability of the recommendation process. Visualization leverages visual representations to facilitate human perception, while interaction stresses user involvement through dialogue with the system.

We have presented an interactive visualization to support exploration, transparency and controllability of recommendations at the ACM Conference on Intelligent User Interfaces (IUI) in 2013 (Verbert et al., 2013). Several other researchers have proposed interactive visualizations as a means to support interaction with recommender systems. In this article, we analyze these interactive recommender systems and their support to address the following challenges: (1) transparency and justification, (2) user control over the recommender system, (3) lack of diversity, (4) cold start issues and (5) contextual information acquisition and representation. The research contributions are

*Corresponding author. Tel.: +32 16 32 82 86; Fax: +32 16 32 79 96.

E-mail address: chen.he@cs.kuleuven.be

Email addresses: chen.he@cs.kuleuven.be (Chen He), dparra@ing.puc.cl (Denis Parra), katrien.verbert@cs.kuleuven.be (Katrien Verbert)

three-fold:

1. We present an interactive visualization framework for recommender systems. The framework integrates visualization and recommendation techniques to address several issues of recommender systems, including cold start, user control and transparency.
2. Then we present an analysis of existing interactive recommender systems along the dimensions of our framework.
3. Based on the analysis, we identify future research challenges and opportunities to advance the research field.

The article is organized as follows: we present recommendation algorithms and visualization techniques in Section 2. Section 3 presents our interactive visualization framework for recommender systems and elaborates the research objectives of our work. Section 4 presents a comprehensive overview of existing interactive recommender systems. An analysis of these systems along the dimensions of our framework is elaborated in Section 5. Finally, we present future research directions and challenges based on our analysis.

2. Background

2.1. Recommendation algorithms and their limitations

Recommender algorithms are often broadly categorized in three areas: *collaborative filtering* recognizes commonalities between users or between items on the basis of explicit (ratings, tags, etc.) or implicit (actions like reading, downloading,) relevance indications (Burke, 2010). A standard user-based collaborative filtering algorithm first identifies similar users based on their overlapping interactions or similar ratings of common items. It then makes recommendations based on preferences of these similar users. A standard item-based recommendation algorithm analyzes similarities between items and then uses these similar items to identify the set of items to be recommended. Collaborative filtering is the most widely implemented and most mature technology (Burke, 2002). *Content-based filtering* matches descriptions of items to descriptions of users (Pazzani and Billsus, 2007). They base their predictions on information about individual users and items, and ignore contributions from other users. This approach relates most closely to our work on metadata (Ternier et al., 2009). *Hybrid recommender systems* combine recommendation techniques, to gain better performance with fewer drawbacks (Burke, 2002).

Recent research on recommender systems is increasingly oriented towards incorporation of contextual information into the recommendation process (Adomavicius and Tuzhilin, 2005). While traditional recommender systems represent the users with simple user models, context-aware recommender systems consider additional information to improve quality of recommendations. For instance, a movie recommender based on collaborative filtering represents the user as a vector of ratings over a set of films, but a context-aware recommender can consider who is accompanying the user – a child or another adult –

to make a more appropriate suggestion. Although “user company” is a typical example of context, there is no clear consensus about its definition and several disciplines understand context differently (Adomavicius and Tuzhilin, 2012). Despite this, a well cited definition of Dey (2001) states that context is “*any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves.*”

Dourish (2004) expands this definition by considering context from social and technological perspectives (Adomavicius and Tuzhilin, 2012). The social perspective understands context as something describing interactions rather than a setting or situation. The technical perspective represents context as a predefined set of observable attributes. We have analyzed several technical definitions of context in Verbert et al. (2012). In summary, most definitions include attributes to represent location, time, computing context, user context, activity of the user, physical conditions such as weather and noise level, or social relations. Context-aware recommender systems try to adapt recommendations to one or more of these contextual attributes and have been proven to provide better predictive performance in a number of domains. Emotion is one of the most popular contextual attributes (Zheng et al., 2013). Examples of other attributes that have been considered in context-aware recommender systems include weather (Hong et al., 2009) and noise level (Yau and Joy, 2007).

Contextual information can be obtained in a number of ways, including *explicitly* from the user or *implicitly* from the environment, for instance by obtaining the current location or device type. Contextual information can also be *inferred* by analyzing user interactions with tools and resources, for instance to estimate the current interest of the user.

Although these algorithms have been implemented and validated on a large scale in several application areas (Nageswara and Talwar, 2008), there are important challenges that need to be addressed before recommender systems can realize their full potential:

1. Collaborative recommendation techniques often suffer from *cold start* issues, i.e. they cannot make effective recommendations for new users or for new items that have no explicit or implicit relevance indicators yet (Burke, 2010).
2. It is *difficult to explain* the rationale behind recommendations to end users (Herlocker et al., 2000): the complexity of recommendation algorithms often prevents users from comprehending recommended results and can lead to *trust issues* when recommendations fail. This complexity is often aggravated by contextual recommendation algorithms that use various types of contextual information in the recommendation process.
3. *Contextual information* can be substantially enriched in non-obtrusive way by exploiting new sensors, particularly in mobile devices like smart phones or tablet computers. In addition, there is a need for developing *richer interaction capabilities* for contextual recommender systems (Adomavicius and Tuzhilin, 2012). The current black box

nature of recommender systems prevents users to provide input into the recommendation process in an interactive and iterative manner. As the predication of the current task or interest of the user is a challenging task, there is a need to develop mixed approaches that enable users to help steer this process.

2.2. Visualization techniques

Data visualization is a well established research field. The distinction is often made between *information visualization* and *scientific visualization*. Information visualization focuses on representing abstract data. A typical example is a graph visualization that shows relationships between people or a time line visualization that represents the evolution of concepts over time. Scientific visualization is specifically concerned with data that has a well-defined representation in 2D or 3D space. Emphasis is on realistic renderings of volumes, surfaces, illumination sources, etc.

In this article, we are most interested in *information visualization*: the use of interactive visual representations of abstract data to amplify cognition (Card et al., 1999). This approach is increasingly applied in scientific research, digital libraries, data mining, financial data analysis, market studies, drug discovery, etc. (Shneiderman and Bederson, 2003).

Research on information visualization is focused on enabling users to control the process of navigating through information spaces in flexible ways. Whereas recommendation algorithms find interesting items in large data sets automatically, information visualization makes use of the principles in Gestalt Theory that explain the human visual capacity, such as proximity, similarity, continuity, symmetry, closure and relative size (Ware, 2000). These principles explain how users see patterns in data.

Information visualization relies on the design of effective and efficient interactive visual representations that users can manipulate to solve specific tasks. This approach is especially useful when a person does not know what questions to ask about the data or when she wants to ask better, more meaningful questions (Fekete et al., 2008). Especially relevant is the intersection of information visualization and search interfaces, where rich results can provide exploration, insight and understanding (Morville, 2005; Ahn and Brusilovsky, 2009). Several data type taxonomies have been described in literature (Chi, 2000; Keim, 2002; Adnan et al., 2008; Ellis and Dix, 2007). For each data type, appropriate visualization techniques and visualization tasks have been designed (Shneiderman, 1996), including:

- histograms, word clouds and box plots (Willett et al., 2007) for 1-dimensional data;
- scatter plots, matrices, linked histograms etc. for 2-dimensional data;
- 3D scatter plots or metaphoric worlds (Santos et al., 2000) for 3-dimensional data;
- timeline visualizations such as theme rivers (Nowell et al., 2002), clustered time series (Van Wijk and van Seelow, 1999) or time matrices (Yi et al., 2010);

- stacked displays such as tree-maps (Shneiderman and Johnson, 1991), sunbursts (Stasko and Zhang, 2000), hyperbolic trees (Lamping and Rao, 1996), dendograms, cone and radial trees (Nussbaumer, 2005) for hierarchical data;
- node-link diagrams (Elmqvist and Fekete, 2010) with graph layout algorithms such as Reingold and Tilford, H-trees and Balloon graphs (Herman et al., 2000) for representing relationships. Venn diagrams, Euler diagrams and cluster maps (Verbert et al., 2013) are used for representing relationships between sets;
- elastic lists (Stefaner et al., 2008), parallel coordinates (Inselberg, 1985), data meadows (Elmqvist et al., 2008), etc. for multi-dimensional data.

This taxonomy has been widely accepted and has been extended with interaction technique taxonomies that consider interactive filtering, zooming, distortion, linking and brushing, etc., as well as task taxonomies for visualization interfaces such as overview, zoom, filter, details-on-demand, relate, history and extract (Keim, 2002).

In this article, we rely on these taxonomies to analyze the visualization and interaction techniques that are used for interacting with recommender systems.

3. Interactive recommender framework

We propose a tight integration of visualization and recommendation techniques to enable end users to interact with recommender systems and to create a feedback loop. The framework shown in Figure 1 explains such an integrated visual recommendation process and the feedback loop that incorporates user feedback and input. The *user data* node refers to user ratings, browsing / search history, etc., which is used as a basis for calculating personalized recommendations. Contextual recommender systems incorporate contextual information for generating recommendations tailored to the current needs of the user, such as location, current activity or interest of the user. Such information is denoted by the *context* node. The *recommender engine* node gets the information from the *user data* node and the *context* node to calculate the data for the *medium* node and the *recommendations* node. The *medium* node represents data inferred from user data and context data by recommender engine: a list of users that are similar to the active user is a typical example of such data. In this particular example, this data is used in a next step by a collaborative filtering recommender engine to generate recommendations based on interests of these like-minded users. By visualizing these similar users, the user is provided with insight of the reasoning behind the recommendations. These recommendations are represented with the *recommendations* node.

The recommendation process is illustrated by the arrows in Figure 1. The straight arrows indicate the data flow while the revolving arrows refer to user interactions with data elements of the different nodes. For instance, the revolving arrow of




Figure 1: Interactive recommender framework

the user data and context nodes represent interaction of end users with a visualization that represents user data and context data, respectively. The revolving arrow of the recommendations node represents interaction of end users with a visual representation of recommendations. Likewise, the revolving arrow of the medium node represents interaction of end users with a visualization of medium data, such as a list of like-minded users. User feedback through the four nodes is transmitted to the recommender engine through the straight arrows pointing towards the recommender engine node. Then the engine recalculates and transmits the revised data to the medium and recommendations nodes to visualize.

From the user perspective, the *recommender engine* node is hidden whereas the other four nodes are visualized. The visualization represents the whole process of recommendations and involves user interaction to get user feedback. Figure 2 represents the user mental model (Norman, 2002) of our framework. The straight arrows represent automatic data calculations and transformations between nodes, whereas the revolving arrows indicate user interaction as explained above.

Note that the active user who interacts with the recommender system is important in this process. She may interact with any node in the model presented in Figure 2. As we will see, some systems represent the active user explicitly in the visualization to help the end user interpret her relation with the different nodes.




Figure 2: User mental model of interactive recommender systems

Existing interactive recommender systems focus on interactive visualization of different parts of this model. Some systems focus on visualization of just one node, whereas other systems cover multiple nodes and focus specifically on visualizing relationships between the different nodes. The major objective of these visualizations is to address limitations of current recommender systems. More specifically, the use of interactive visualizations is researched to achieve the following objectives:

- *Transparency* deals with the “black-box” nature of current recommender systems by explaining the inner logic of the

system to end users (Sinha and Swearingen, 2002; Vig and Riedl, 2009). For example, visual representations of the neighborhood structure and interests of like-minded users can convey information about interests of peers (Gretarsson et al., 2010; Klerkx and Duval, 2009) and help users to identify how and whether interests of users in their neighborhood match their own interests or needs (O’Donovan et al., 2008). User understanding of the reasoning behind a recommendation may help to increase confidence in that recommendation (Herlocker et al., 2000; Abdul-Rahman and Hailes, 2000).

- Similar to transparency, *justification* helps users understand why they get certain recommendations, but it may not relate to the inner logic of the recommendation techniques (Vig and Riedl, 2009; Tintarev and Masthoff, 2011). That is, if the system only describes why the user gets the recommendations and does not describe how the recommendation is selected or how the system works, then it only justifies the recommendations (Tintarev and Masthoff, 2011). For example, Amazon.com explains why the items are recommended by mentioning ‘we recommend these products based on the products you recently purchased’. Bogdanov et al. (2013) justifies the recommendation by mapping the description of user preferences to graphic symbols, such as a guitar representing rock music. Both approaches explain recommendations, but do not provide insight into the recommendation techniques that are used. In some circumstances, justification may be more preferred than transparency. For example, the recommendation technique may be too complex to describe or designers intend to keep the inner logic hidden (Vig and Riedl, 2009; Tintarev and Masthoff, 2011; Herlocker et al., 2000).
 - *Controllability* strengthens user involvement by incorporating input and feedback from the end user into the recommendation process. User control can occur at any step of the recommendation process, such as providing ratings, adjusting preference data, and revising or exploring recommendations. As an example, TasteWeights (Bostandjiev et al., 2012) allows users to fine-tune the weights on different parameters to customize recommendations.
- Understanding the relationship between the input and output of the system can enable the user to meaningfully revise input parameters and thus improve recommendations (Swearingen and Sinha, 2001). It is useful to compensate for deficiencies in recommendation algorithms and allows users to tailor recommendations to their rapid changing preferences. The intent of proper user control is to increase recommendations accuracy (Pu et al., 2012) by incorporating user input and feedback. In general, previous work shows a positive relationship between user satisfaction and user control (Parra and Brusilovsky, 2015).
- *Diversity* refers to providing recommendations with a relatively large coverage of the recommendation space (Hu and Pu, 2011). For instance, it is important to recommend

items that the user would prefer, but that are different from those which she has already purchased or experienced. Related work shows that recommendations should maintain a certain level of diversity, even if it sacrifices overall accuracy (Pu et al., 2012). However, this research also shows that predicted diversity is not directly correlated with perceived diversity, so there is a need to leverage visualization design to enhance the perceived diversity in recommendations. For example, Hu and Pu (2011) show that visualizing recommendations in categories rather than a list enhances user perception of the recommendation diversity and has a positive effect on acceptance of recommendations.

- When a new item or a new user joins a recommender system, the system has no prior knowledge about it, i.e., no item-feature data, no ratings, no preference information. The inability to make recommendation to new comers is called the *cold start* problem (Schein et al., 2002). This problem can be alleviated by algorithmic approaches, for instance by clustering particular items or users (Halder et al., 2012). Also conversational recommendation interfaces (Felfernig and Gula, 2006) have been introduced that elicit user preferences in a way that reduces perceived efforts of users. In this article, we focus on interactive visualization techniques to tackle this challenge. An example is a visual overview of popular content to enable new users to locate their interests in a straightforward way (Zhao et al., 2010).
- Acquiring *contextual* information and incorporating it into recommendation processes in a flexible and fluid manner has gained increased interest over the past decades. The goal is to tailor recommendations to the current needs of the user. Among various contextual data, *emotion* is an important contextual criterion and plays a key role in decision making (Picard et al., 2004). Lerner et al. (2015) summarize eight themes that stress this strong impact on decision making. Integrating affective elements in recommender systems is challenging, among others because estimation of such variables in an automatic way is difficult. In this article, we focus again on solutions that use visualization techniques to capture user input and the role that these visualizations play in recommender systems.

4. Survey of interactive recommender systems

In this section, we present a survey of existing interactive recommender systems. The systems are clustered by the objectives that we defined in the previous section, but some systems may address more than one objective.

This list is by no means exhaustive, but it presents nevertheless a broad range of interesting work in this area. We analyze the commonalities and differences of these visual interfaces in the next section.

4.1. Transparency

In total, we surveyed a set of 24 systems that introduce interactive visualizations on top of recommender systems. Eleven out of these 24 systems focus on the use of visualization to support transparency. Several systems explain the process of collaborative filtering. Figure 3 shows the main interface of Peer-Chooser (O'Donovan et al., 2008) for explaining user-based collaborative filtering. Among others, the interface highlights similar users around the active user. The degree of similarity is indicated through their distance to the active user.

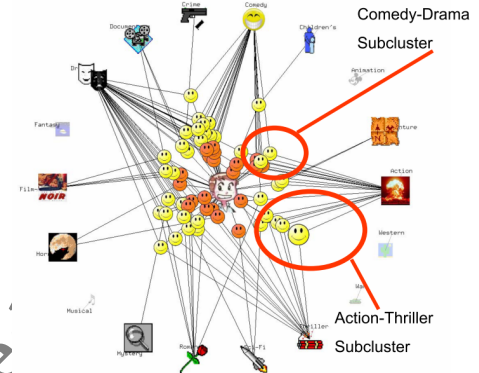


Figure 3: PeerChooser (O'Donovan et al., 2008) visualizes the active user in the center. Similar users are represented around the active user. Recommended items are represented in the outside region. Line connection and distance indicate relationships between nodes and the degree of similarity. [used with permission]

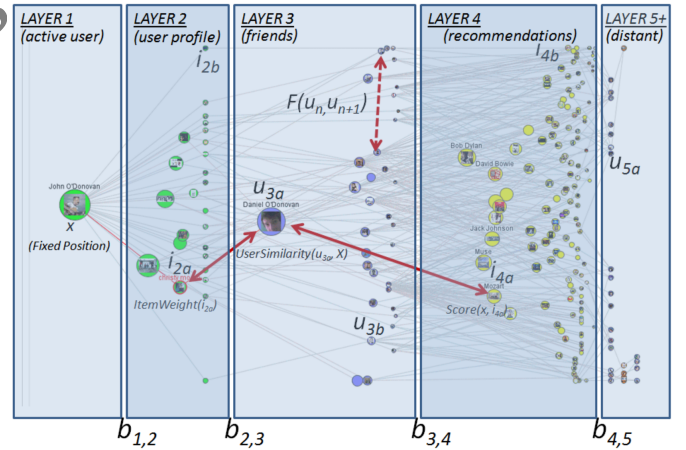


Figure 4: SmallWorlds (Gretarsson et al., 2010) arranges nodes and connections between nodes in five layers: the active user, user profile items, similar users, recommendations, remaining friends. Line connections, size and distance between nodes indicate their relationship. [used with permission]

Similar to PeerChooser, SmallWorlds (Gretarsson et al., 2010) visualizes the inner logic of collaborative filtering recommendations. Five columns are represented: the active user, user profile items, similar friends, recommendations and remaining friends. Information such as item weights and friend similarity are represented by the position and size of the nodes (Figure 4).

TasteWeights (Bostandjiev et al., 2012) generates recommendations with multiple techniques and data sources. The system interconnects user ratings, calculated user preferences, and recommendations to explain the provenance of recommended items (Figure 5). LinkedVis (Bostandjiev et al., 2013) and Schaffer et al. (2015) use the same visualization approach as TasteWeights, but in a different context and with different data sources.

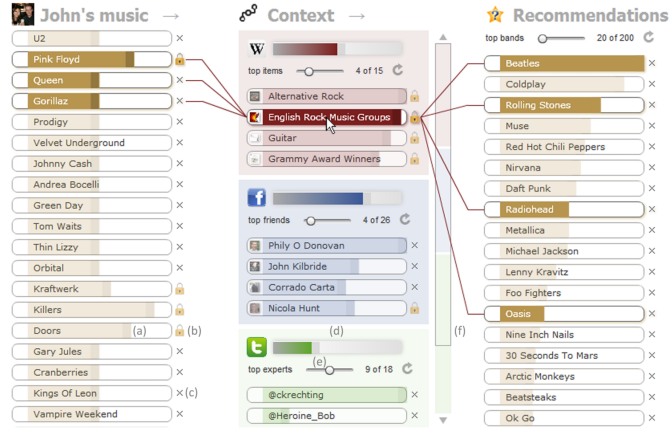


Figure 5: TasteWeights (Bostandjiev et al., 2012) represents the process of recommendation in three connected layers and enables the user at the same time to fine-tune each node. [used with permission]

Graph embeddings (Vlachos and Svonava, 2012) and TIGRS (Bruns et al., 2015) visualize a collection of similar items to a pivot item. Similar to PeerChooser and SmallWorlds, they use a node-link diagram and distance between nodes to indicate their similarity. Color cues are used to cluster items of the same type in Graph embeddings (Figure 6). TIGRS represents recommendations and links to related keywords that match the user interest as a basis to explain the recommended items (Figure 7).

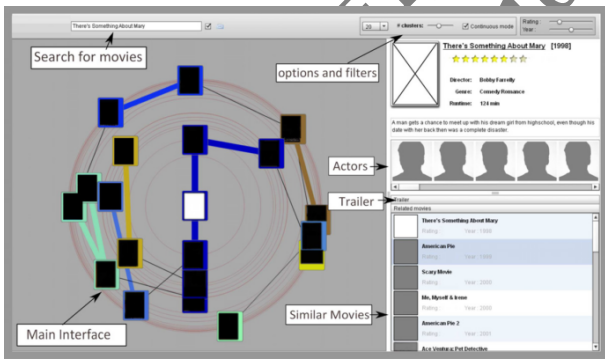


Figure 6: Graph embeddings (Vlachos and Svonava, 2012) uses a node-link diagram to visualize similar movies to a pivot movie. Color is used to cluster the same type of movie. Users can input the pivot movie as well as modify the number of clusters and filter items by rating or year using sliders. [used with permission]

TalkExplorer (Verbert et al., 2013) and SetFusion (Parra et al., 2014) visualize relationships between recommendations and multiple recommendation techniques. In TalkExplorer (Figure 8), recommendations of multiple recommendation tech-

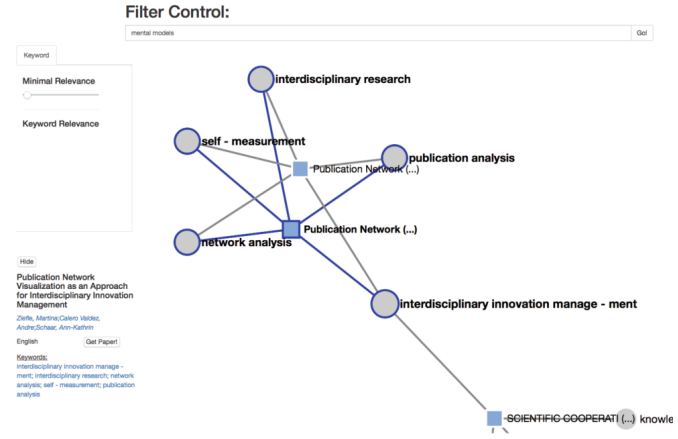


Figure 7: TIGRS (Bruns et al., 2015) visualizes recommended items and related keywords in a node-link diagram and uses filter controls and threshold adjustments of keywords to refine recommendations. [used with permission]

niques are represented as agents, such as a tag-based agent and a content-based agent that use a tag-based and content-based recommendation technique, respectively. Users can browse and interrelate recommendations of these agents, and explore relationships with bookmarks of other users and tags to find relevant items. The system uses a cluster map visualization.

SetFusion (Figure 9) uses a Venn diagram to examine and filter items recommended by multiple techniques. The interface is separated into three parts: the importance of each technique is represented in the top left corner, the Venn diagram in the bottom left corner represents relationships between the recommendations and the techniques, and the recommendation list on the right side represents the details of the recommendations. Color cues are used to connect the three parts. For instance, the color cues next to an item in the recommendation list are consistent with the colors of the recommendation techniques. Similar to TasteWeights and LinkedVis, the approach enables users to understand the inner logic of a hybrid recommender system. That is, a user can see which items are recommended by which recommendation techniques and the importance or weights of each of the techniques in the recommendation process.

Different from previous visualizations, PARIS (Jin et al., 2015) visualizes user characteristics of the user profile (personality traits, age and gender) and the recommendation process, i.e. which information is used and in which order to generate recommendations (left part of Figure 10). Finally, SFViz (Social Friends Visualization) (Gou et al., 2011) uses a Radial Space-Filling (RSF) technique (Chuah, 1998) to visualize a social network hierarchically. Figure 11(a) shows an example: the top 10 recommended friends are represented with colors from red to yellow (highly relevant to less relevant). With edge bundling in Figure 11(b), the rationale of recommendations is provided by showing how the user is connected to the recommended user. In this example, the active user and the recommended user have shared friends in the "hip hop" category.




Figure 8: TalkExplorer (Verbert et al., 2013) uses a cluster map to visualize relations between recommender agents (content-based agent and tag-based agent in the example) and bookmarks of users. Users can explore which items (represented by the yellow bubbles) are recommended by which agents and can examine relationships. For instance, the set of items that is labeled with number 1 represents items that are recommended by both agents. The set of items labeled with number 2 represents items that are recommended by the tag-based agent and are also bookmarked by user “L. Aroyo”. [used with permission]

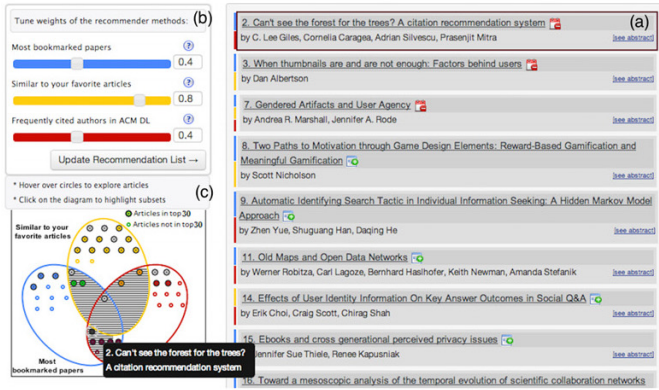


Figure 9: SetFusion (Parra et al., 2014) visualizes relationships among recommended items and multiple recommendation techniques with a Venn diagram and color cues. These color cues are used to represent the different techniques of a hybrid recommender and are used to link the recommendation results to the techniques that produced these recommendations. [used with permission]

4.2. Justification

Justification also helps the user to understand why she gets certain recommendations, but different from transparency this explanation may not relate to the inner logic of the recommendation process.

Seven systems out of the 24 systems that we surveyed are designed to support justification. Tagsplanation explains the recommendations by interrelating two key components: *tag relevance* (the degree to which a tag describes an item) and *tag preference* (the user’s sentiment towards the tag) (Figure 12). Movie Explain provides justification by presenting movie features that the user likes in a table, as illustrated in Figure 13. [The reason is] indicates the user preferred feature, [because you rated] indicates the rating history of the user. Many other text-based approaches to support justification have been surveyed in Tintarev and Masthoff (2011).

Four systems justify recommendations by visualizing user profile data. Bakalov et al. (2013) represent these interests in

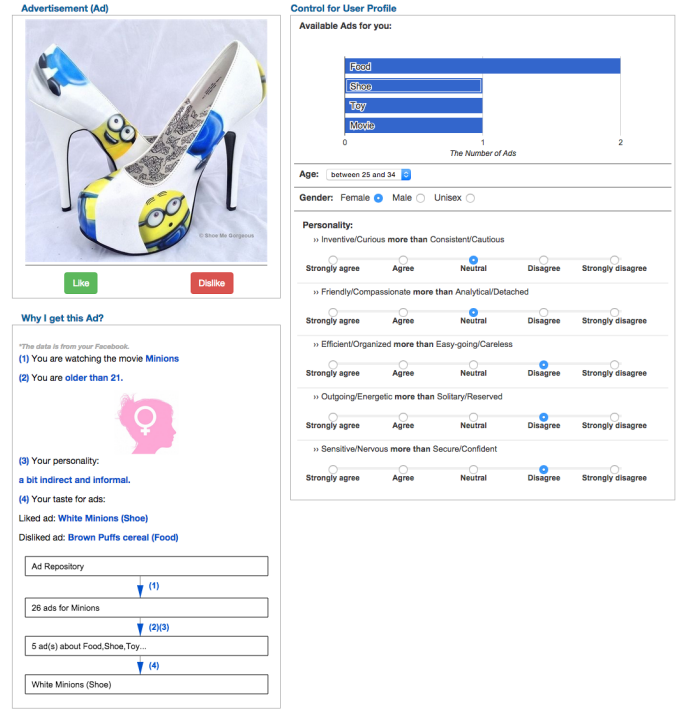


Figure 10: PARIS (Jin et al., 2015) represents user characteristics of the user profile (age, gender, personality) and the recommendation process - i.e. which data is used in which step (bottom left corner). The right side panel supports user control. [used with permission]

zones partitioned into slices, where each zone represents keywords of a certain interest degree, from interesting (center) to uninteresting (Figure 14), and each slice (in between black lines) represents keywords of a specific type. Similar zones are used by work of Kangasrääsio et al. (2015) that is presented in Figure 15. System U (Badenes et al., 2014) represents personality traits using a Sunburst technique to give insight into these variables of the user profile. Bogdanov et al. (2013) maps user preferences to a set of graphical symbols as a means to justify recommendations. Figure 16 shows an example: both the turntable and the blue short hair represent user preferences for electronic and danceable music.

Finally, MusiCube (Saito and Itoh, 2011) integrates user data and recommendations into a single view, illustrated in Figure 17. The system visualizes the distribution of positive (pink) and negative (blue) ratings of the user in the feature space through a scatter plot. Recommended items (yellow dots) are represented in the same scatter plot. The spatial relationship of rated and recommended items facilitates visual perception of their similarity.

4.3. Controllability

Fourteen systems out of the 24 systems that we surveyed support user control. Eleven systems allow *user intervention* into the recommendation process. The remaining three systems support *user exploration* that enables users to navigate through the information space as a means to find other relevant items.

PeerChooser is an example of the first category and enables the user to move a certain genre node closer to the represen-



Figure 11: An example of SFViz (Gou et al., 2011) recommending friends. (a) shows the top 10 recommended friends. The view shows the hierarchical relationship among them. The edge bundling in (b) shows how the active user is connected to the recommended user: in this example because they have shared friends in the “hip-hop” category. [used with permission]




Figure 12: Tagsplanation (Vig and Riedl, 2009) justifies recommendations by representing the relationship between tag relevance (the degree to which a tag describes an item) and tag preference (the user's sentiment towards the tag). [used with permission]

Our Justified Recommendations				
[Movie id]	[Movie Poster]	[Movie title]	[The reason is]	[because you rated]
176		Aliens (1986)	Cameron, James (I)	4 movies with this feature
930		Chain Reaction (1996)	Freeman, Morgan (I)	3 movies with this feature

Figure 13: MoviExplain (Symeonidis et al., 2009) justifies the recommendations in a table that represents relationships between the rating history of the user and a feature of the recommended item. [used with permission]

tation of the active user to increase its weight in calculating recommendations. Similar interactions are supported by Small-Worlds and are illustrated in Figure 4. In work of Bakalov et al. (2013) and Kangasrääsiö et al. (2015), the user can adjust her profile by dragging a keyword over the circular layout to change its interest level.




Figure 16: Musical avatar (Bogdanov et al., 2013) justifies recommendations by representing the user profile using graphical symbols. [used with permission]




Figure 17: MusiCube (Saito and Itoh, 2011) visualizes positive (pink) and negative (blue) ratings of the user. Recommended items (yellow dots) are represented in the same scatter plot and enable the user to explore relationships between these recommended items and her previously rated items. [used with permission]

threshold of relevance for each keyword, illustrated in the left part of Figure 7.

PARIS allows the user to adjust her profile with input controls such as drop-down lists and check lists, illustrated in the right part of Figure 10. MusiCube allows user intervention by enabling the user to rate more items to refine the recommendations directly in the visualization - i.e. by selecting a yellow dot that represents a recommended item.

Three systems support *exploration* of the recommendation space to find other relevant items. Graph embeddings allows users to input a pivot item as a basis to find other relevant items. Users can also modify the number of displayed clusters and filter items by ratings and publication year, illustrated in Figure 6. Similarly, in SFViz the user can specify a category of interest.

TalkExplorer, presented in Figure 8, enables users to explore and combine multiple recommendation techniques, users and tags. These entities can be added to the visualization as a basis to support exploration.

4.4. Diversity

To the best of our knowledge, only one system focuses on visualizing diversity of recommendations. The Diversity Donut (Wong et al., 2011) is an interactive recommender system that allows a user to control the level of opinion diversity by shrinking the donut to see responses from like-minded users, or expanding the donut to see responses from users who differ in opinion. The approach is illustrated in Figure 18 and enables users to adjust the coverage of recommendations.




Figure 18: The Diversity Donut (Wong et al., 2011) visualizes item diversity in different levels. Distance to the center is used to represent the degree of similarity, with more similar items closer to the center. [used with permission]

4.5. Cold start

Three systems out of the 24 systems that we surveyed explore ways to alleviate the cold start problem. Pharos (Zhao et al., 2010) addresses the cold start problem by providing an overview of popular communities on a website. Figure 19 shows an example of this social map visualization. A community is identified by a set of people (blue) and content (green). The size of the text indicates their importance. Inactive items are represented in gray. Novice users can interact with this visualization to locate their interests.



Figure 19: Pharos (Zhao et al., 2010) recommends popular sets of items for novice users by clustering similar content (green) and people (blue) into communities. The position close to the center and size of the community indicates the popularity. [used with permission]

MrTaggy (Kammerer et al., 2009) and Loepp et al. (2014) activate recommendations through dialog with the user and filter them in an iterative manner. In a first step, MrTaggy uses user selected or inputted keywords. The system then recommends related keywords and allows the user to specify the relevance of these related keywords with upward and downward arrows, as illustrated in the left part of Figure 20. The system also uses such controls to elicit user feedback on the relevance of recommendations. Loepp et al. (2014) elicits user preferences by asking the user to choose iteratively between two sets of sample items that represent low and high values of a certain factor, respectively (Figure 21). Each interaction step contributes to a more precise positioning of the user in the feature space.

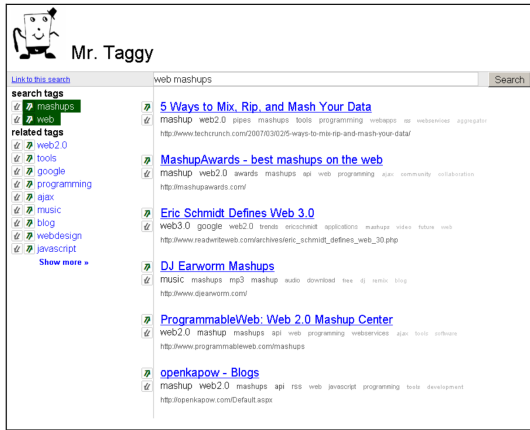


Figure 20: MrTaggy (Kammerer et al., 2009) recommends tags on the left part of the interface and shows recommendations on the right part. The system enables user control of the relevancy of both tags and recommendations through upward and downward arrows. [used with permission]

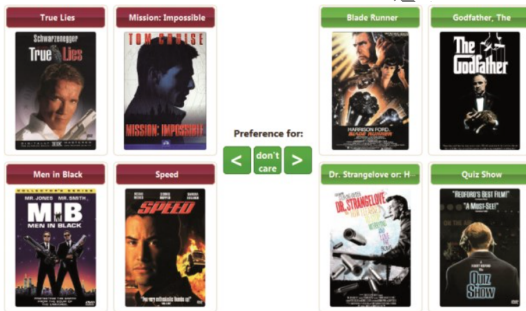


Figure 21: Loepp et al. (2014) allows the user to choose iteratively between two sets of sample items that represent low and high values of a certain factor respectively to elicit user preferences [used with permission].

4.6. Context

Two systems out of the 24 systems that we surveyed focus on using interactive visualizations to incorporate contextual information, and more specifically emotions of the user, into the recommendation process. As described above, emotional criteria are key in decision making and play an important role in recent research on recommender systems.




Figure 22: CoFeel (Chen and Pu, 2014) represents emotions with different colors on a plate. [used with permission]

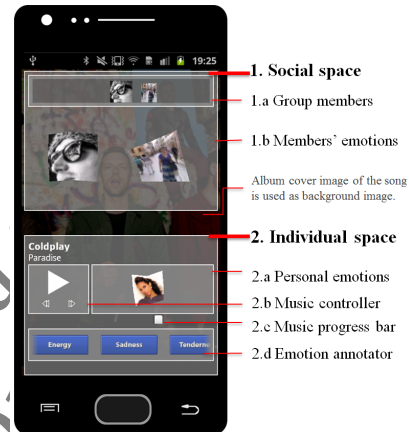


Figure 23: Empatheticons (Chen et al., 2014) uses deformations of user profile pictures to show emotions. Its implementation allows the user to view emotion feedback of other users and control her own feedback. [used with permission]

CoFeel (Chen and Pu, 2014) and Empatheticons (Chen et al., 2014) focus on explicit emotional input and feedback visualizations in group recommender systems and are illustrated in Figure 22 and Figure 23. CoFeel is designed as an emotional plate based on a Geneva Emotion Wheel (Scherer, 2005). It uses the plate-hole-ball metaphor to elicit user input, i.e. users can select the emotion by placing the ball on a certain emotion. Empatheticons are a set of animated icons to represent different emotions. For example, the animation of the emotion *joyful* utilizes the metaphor “leaving the ground and up in the air”. CoFeel and Empatheticons are both incorporated into a group music recommender system as emotional input and visualization methods. Users can provide feedback to a recommended item through these interfaces and see emotional feedback of other users.

5. Analysis

In this section, we analyze the systems presented in the previous section. We present an analysis along the dimensions of our framework presented in Section 3. The analysis results are again clustered by the main objectives of the systems. In addition, we analyze the visualization techniques and recommendation algorithms that are used, and evaluation results that assess

	Visualization Objectives						Visualization Techniques						Recommender Algorithms				
	Transparency	Justification	Controllability	Diversity	Cold start	Context	Node-link diagram	Set-based view	Radial view	Table	Scatterplot	Icon	Flow chart	Collaborative filtering	Content-based	Hybrid	Group Recommender
Bakalov et al. (Figure 14)		+	+						+						+		
Bogdanov et al. (Figure 16)		+										+			+		
CoFeel (Figure 22)						+			+								+
Diversity Donut (Figure 18)				+					+					+			
Empatheticons (Figure 23)						+						+					+
Graph embeddings (Figure 6)	+		+				+								+		
Kangasrääsiö et al. (Figure 15)		+	+						+						+		
LinkedVis	+		+				+							+			
Loepp et al. (Figure 21)					+			+						+			
MoviExplain (Figure 13)		+								+				+	+		
MrTaggy (Figure 20)					+							+			+		
MusiCube (Figure 17)		+	+								+				+		
PARIS (Figure 10)	+		+										+		+		
PeerChooser (Figure 3)	+		+				+							+			
Pharos (Figure 19)					+			+						+			
Schaffer et al.	+		+				+							+			
SetFusion (Figure 9)	+		+					+								+	
SFViz (Figure 11)	+		+						+					+			
SmallWorlds (Figure 4)	+		+				+							+			
System U		+							+							+	
Tagsplanation (Figure 12)		+								+				+			
TalkExplorer (Figure 8)	+		+					+								+	
TasteWeights (Figure 5)	+		+				+									+	
TIGRS (Figure 7)	+		+				+								+		

Table 1: Analysis of interactive recommender systems

the impact on the recommendation process. An overview of the analysis is shown in Table 1 and Table 2. These tables include references to the figures of the different visualizations.

5.1. Objectives

5.1.1. Transparency

Transparency is supported by 11 systems of our survey. Figure 24 represents which nodes of our framework are visualized by the surveyed systems to explain the inner logic of a recommender system. PARIS visualizes the medium node and uses a sequence graph to show which data of the user profile is used in different steps of the recommendation process.

Five systems visualize the relationship between the medium and the recommendations. For instance, PeerChooser shows relations between recommendations and similar users (medium) as a means to explain collaborative filtering results.

In addition to visualizing the relationship between the medium and recommendations, five systems represent relationships with user data. SmallWorlds represents for instance relationships between user preferences (user data), similar users (medium) and recommendations through line connections.

5.1.2. Justification

Seven systems of our survey support justification through visualizations. The systems also explain recommendations, but

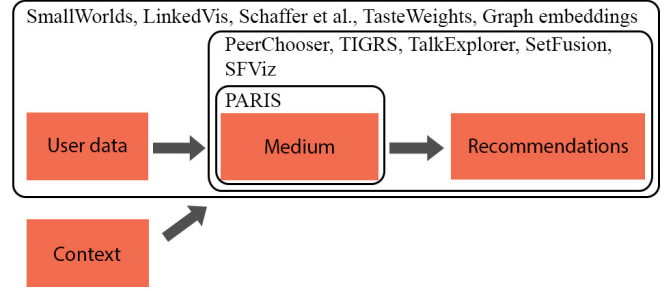


Figure 24: Systems visualize *transparency* in different ways. PARIS focuses on visualizing the medium node. Five visualizations also depict the relation with recommendations. Five systems visualize relationships among user data, medium and recommendations.

do not provide insight into the inner logic of the underlying algorithm. Four visualizations give the user insight into the medium node as a basis to explain recommendations, as illustrated in Figure 25. Bogdanov et al. (2013) for instance justify recommendations by representing the user profile with avatars. Although user profile values are represented to the user, the system does not explain how this information is used to generate recommendations. PARIS is an example that does provide such insight into the recommendation process: i.e. the system also shows how this information is used and in which order as a basis to support not just justification, but also transparency of the inner logic.

Two systems justify recommendations by representing the relationship between a recommendation and the value of a specific attribute of the medium node. Tagsplanation for instance represents tags and the user sentiment towards that tag (medium) to explain recommendations. In addition, MusiCube presents the relation with user data to justify recommendations. The approach enables users to see correlations between their ratings and recommendations.

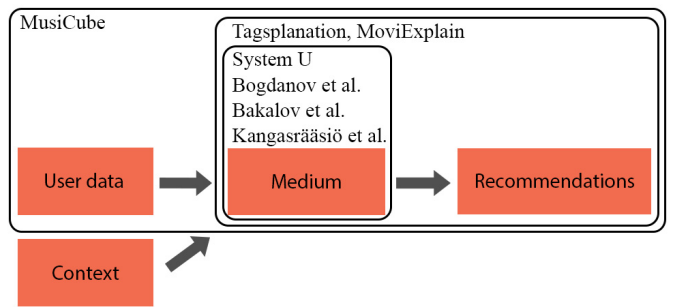


Figure 25: Four systems *justify* recommendations by representing the medium node. Two systems also present relations with recommendations. MusiCube depicts the relationships among the three nodes: user ratings, music features and recommendations.

5.1.3. Controllability

Fourteen systems of our survey support user control to enable the user to intervene in the recommendation process as a basis to improve recommendations or to explore the recommendation space, as illustrated in Figure 26.

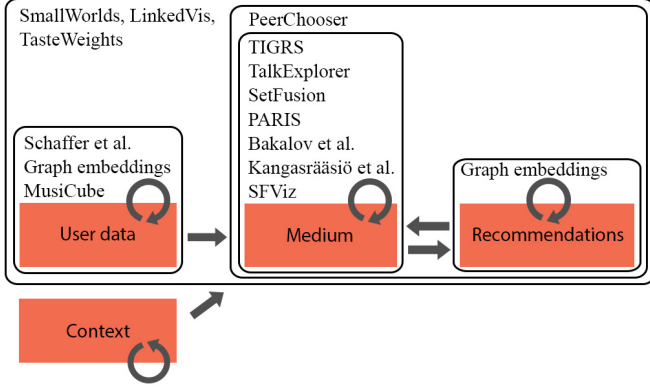


Figure 26: Three systems focus on *control* over user data. Seven systems represent medium data and one system focuses on filtering the recommendations. PeerChooser allows control over both medium and recommendations. Three systems enable user control over the three nodes.

Three system of our survey focus on input (Vlachos and Svonava, 2012; Saito and Itoh, 2011) or adjustment (Schaffer et al., 2015) of user data. These systems use visualization techniques to elicit additional or adjusted user input such as ratings in a straightforward way as a basis to improve recommendations. Graph embeddings also focuses on interactive visualizations of the output (recommendations) as a basis to refine the recommendation list.

Seven systems enable user control over the medium node. For example, PARIS and Bakalov et al. (2013) enable the user to adjust her user profile. In PARIS, the user profile represents user personality characteristics of the user model inferred from user data of Facebook. In work of Bakalov et al. (2013), the user profile represents keywords the user is interested in derived from queried publications. The user can adjust these user profile values to improve recommendations. PeerChooser also visualizes relationships to recommendations. Users can adjust the position of similar users to update recommendations.

Three systems involve the user in all three stages of the recommendation process. For instance, SmallWorlds enables the user to adjust the position of the user preferences (user data), similar user items (medium) and recommended items to refine recommendations.

5.1.4. Diversity

The Diversity Donut visualizes the different levels of recommendation diversity. The system uses a circular layout to depict the diversity level of each recommended item through its distance to the center, illustrated in Figure 27.

5.1.5. Cold start

In our collection, the cold start problem is addressed in two ways. The first approach, used by Pharos, recommends the most popular content to novice users and hence visualizes the contents of the recommendation node, as illustrated in Figure 28.

The second approach uses a drill down technique in an iterative manner to elicit user interests. Loepp et al. (2014) use this




Figure 27: *Diversity* is supported by one system which focuses on visualizing the recommendations node.

approach on user data: the system enables user selection between two sets of different items to incrementally build the user profile. MrTaggy uses this approach on medium data. The system first suggests related keywords (medium) to a user selected or inputted keyword. These related keywords are then used to narrow down recommendations.




Figure 28: Three systems alleviate the *cold start* problem in a different way.

5.1.6. Context

Two systems of our survey involve contextual information in the recommendation process. As illustrated in Figure 29, both systems focus on the input level of contextual information. CoFeel for instance uses an emotion plate with different colors to represent different emotions. The visualization of such contextual information is used to elicit input and feedback from end users in an intuitive way.




Figure 29: Two systems focus on the input and feedback of *contextual information*.

5.2. Visualization techniques

Visualization techniques that are adopted by the systems that we have surveyed can be categorized in seven clusters, as illustrated in the middle part of Table 1.

1. Seven systems use *node-link diagrams* to represent relationships. The nodes represent user data, medium content and/or recommendations, and are often clustered into

layers. Then, line connections are used to connect items in the user data, medium and recommendations to explain the provenance of recommended items and to enable user control. Multiple connected layers are represented in a circular way (PeerChooser) or with columns (SmallWorlds, LinkedVis, TasteWeights).

2. *Set-based visualizations*, used by four systems, gather similar items into sets, which facilitates explanation of their commonalities. Compared to line connections in node-link diagrams, TalkExplorer and SetFusion represent recommended items into sets using a cluster map visualization and a Venn Diagram, respectively, to decrease the visual complexity of explanations. Pharos represents popular communities of users and content in clusters to help users locate their interests. Loepp et al. (2014) allows the user to choose from a set of similar items rather than a single item to identify her interests.
3. *Radial visualizations* are used by six systems. SFViz uses a Radial Space-Filling technique to represent a social network, where nested circles indicate parent-child relationships. The approach is less intuitive than traditional node-link diagrams, but can scale more easily to large networks. Similarly, System U adopts a Sunburst technique to visualize hierarchical personality data. The other systems use a radial visualization to represent different levels of a particular variable. These variables can be user interests (Bakalov et al., 2013; Kangasrääsiö et al., 2015), diversity (Wong et al., 2011), or emotions (Chen and Pu, 2013). Bakalov et al. (2013) for instance use the approach to represent user interests by keywords, with those keywords that are of higher interest closer to the center.
4. *Tables* are used by two systems. Both MovieXplain and Tagsplanation use the approach to represent relationships between recommendations and variables that explain these recommendations.
5. MusiCube uses a *scatter plot* that represents ratings of users and recommendations by colored dots. The approach enables users to identify correlations between recommendations and their ratings.
6. Three systems use *icons*. Empatheticons uses animated icons to represent different emotions. Bogdanov et al. (2013) map descriptions of user preferences to graphic symbols to give a user insight into her profile data. MrTaggy uses icons to enable users to indicate the relevance of recommended items and tags.
7. Finally, PARIS uses a flow chart to indicate which information of the user profile is used in which order to generate recommendations.

Overall, most interactive recommender systems use visualizations to represent relationships among data elements. Node-link diagrams (7) are most often used. Other representations for relationship data include set-based visualizations (4), radial visualizations (6), flow charts (1) and tables (2). MusiCube uses an interesting scatter plot approach to enable users to find correlations between ratings and recommendations. The other sys-

tems use icons to represent relevant information. We discuss these techniques in Section 6.5.

5.3. Recommendation algorithms

Collaborative filtering (CF) is used by ten out of 24 systems. There are two main approaches:

1. Four systems, PeerChooser, SmallWorlds, LinkedVis and the work of Schaffer et al. (2015), support transparency of the CF algorithm and all use node-link diagrams that visualize among others *similar users* to explain the rationale of the CF algorithm. SmallWorlds also visualizes dissimilar friends to increase transparency and to support user control.
2. The remaining six systems use different representations to visualize *user data*, *medium data* or *recommendations*. SFViz visualizes similar users with a Radial Space-Filling. Tagsplanation represents the user's sentiment towards relevant tags in a table. The Diversity Donut represents recommendations of a CF technique in a radial view.

Ten systems use *content-based recommendations*. There are three main approaches:

1. Four systems visualize *relationships* among recommendations. Graph embeddings for instance represents similarity on a 2D canvas: distance is used to represent the level of similarity. This visualization is used to enable users to find other relevant items that are suggested by a content-based recommender system. MovieXplain uses both CF and content-based techniques: its interface describes the content features and connections to user rated items to justify recommendations.
2. Five systems represent (parts of) the *user profile* that is built by a content-based recommendation technique. Bogdanov et al. (2013) for instance use icons to represent preferences of the user profile. Bakalov et al. and Kangasrääsiö et al. represent user interest on a concentric layout. System U and PARIS represent among others personality traits of the user profile.
3. Finally, TIGRS visualizes the *important keywords* to each recommended document through a node-link diagram.

Three systems are implemented with *hybrid recommendation techniques*. Their visualizations facilitate identifying which technique is used to generate recommendations, or which combinations of techniques, and users can control the importance of these techniques to tailor recommendations. TasteWeights for instance leverages among others content-based and collaborative techniques and visualizes user preferences from these perspectives in groups. Finally, CoFeel and Empatheticons implement a *group recommender system* and focus on visualizing emotions of group members as a basis to tailor recommendations.

	Approaches					Metrics							Data collection methods					Results							
	Comparison with baseline without recommendations	Comparison with baseline without user control or visual explanation	Comparing different visualizations	Comparing different recommender algorithms	Asking users to explore freely	Effectiveness	Efficiency	Engagement	Satisfaction	Trust	Usability	Usefulness	Recommendation accuracy test	Task performance analysis	User behavior	Questionnaire	Interview	Think-aloud	Increase of acceptance	Better task performance	Increase of efficiency	Increase of engagement	Increase of satisfaction	Increase of trust	Positive usefulness / usability feedback
Bakalov et al. (Figure 14)		+							+	+	+	+				+							+	+/-	+
Bogdanov et al. (Figure 16)			+	+								+				+									+
CoFeel (Figure 22)					+			+			+	+			+		+					+			+
Diversity Donut (Figure 18)				+		+		+								+			-			+			+
Empatheticons (Figure 23)					+			+			+	+			+	+		+			+				+
Kangasrääsiö et al. (Figure 15)		+				+			+	+	+	+		+		+	+			+/-		+/-	+/-	+/-	+/-
LinkedVis		+		+		+					+	+	+			+			+						+
Loepp et al. (Figure 21)	+					+				+		+				+			+				+		+
MoviExplain (Figure 13)			+			+							+			+			+					+	+
MrTaggy (Figure 20)	+					+		+			+	+		+	+	+				+		+			+
MusiCube (Figure 17)					+	+					+		+	+		+			+	+/-					+
PARIS (Figure 10)		+				+			+	+		+			+	+			+			+/-	-		+
PeerChooser (Figure 3)	+	+				+						+	+			+			+						+
Pharos (Figure 19)	+					+	+					+		+		+				+	+				+
Schaffer et al.		+				+			+	+			+	+	+	+			+			+/-	+/-		
SetFusion (Figure 9)		+				+	+	+				+		+	+	+			+	+	+	+			+
SmallWorlds (Figure 4)	+		+			+			+		+	+	+			+			+			+			+
Tagsplanation (Figure 12)			+			+						+				+			+/-						
TalkExplorer (Figure 8)					+	+					+	+		+	+	+		+	+	+					+
TasteWeights (Figure 5)		+		+		+					+	+	+			+			+						+
TIGRS (Figure 7)			+			+				+	+					+			+				+		+

Table 2: Evaluations of interactive recommender systems

5.4. Evaluation

To the best of our knowledge, two systems have not been evaluated yet. SFViz exemplifies its visualization framework by use cases.

Of interest for this analysis are the 21 remaining systems that have been evaluated with user studies. There are four main approaches, as presented in Table 2:

1. Thirteen systems have been evaluated by a *comparison with baseline data*. Such baseline data includes data generated by a system without recommendations (5) and data of a system without interaction controls or visual feedback (8).
2. Five systems have been evaluated by *comparing the use of different visualizations* to support the objectives of the system.
3. Four systems have been evaluated by *comparing the use of different recommendation algorithms*.
4. Four systems have been evaluated by *asking users to explore* the system. In these evaluations, all interactions of users are typically logged and analyzed to gain insight into how the system is used and what the effect is on improving recommendations.

These evaluations focus on different evaluation criteria. *Effectiveness* has been evaluated for 17 out of 21 systems. Effectiveness measures whether the interface has an effect on the acceptance of recommendations or task performance. Seven systems assessed the increase of recommendation accuracy with visual explanation or user control through user ratings or automatic accuracy tests. Four systems evaluated the acceptance of recommendations with user subjective feedback. Six systems have assessed impact on task performance.

Whereas generally evaluations show positive results, Loepp et al. (2014) found that manual exploration of movies fits user interests better than interactive recommendations when the user has a specific focus in mind. The authors of PeerChooser found that with user control and dynamic feedback user tends to over-tweak the graph which makes the results over-fitted to specific items. Two studies indicate that there is an inconsistency between recommendation accuracy and perceived accuracy. TasteWeights shows that despite the fact that Wikipedia outperforms Facebook in accuracy, subjects trust recommendations from Facebook more than from Wikipedia. Schaffer et al. (2015) show that users may overvalue their profile updates as perceived accuracy was much higher than the actual accuracy

after user adjustment.

Six systems have assessed the effect on task performance by measuring the quality or productivity of task results. Results of MrTaggy and Pharos indicate an increase of task quality for novice users. Results of SetFusion and TalkExplorer indicate an increase of productivity with transparency and controllability. Kangasrääsiö et al. (2015) show better performance for focused search only, not broad search. For MusiCube the user sample may be too small to draw strong conclusions.

Efficiency and engagement have been evaluated for two and four systems, respectively. Efficiency and engagement compare the time as well as the productivity of performing tasks under different settings. The evaluation of Pharos shows that the system can help users to quickly understand the system and to more efficiently locate their interests compared to a baseline system. User studies of SetFusion show that users are able to find and bookmark recommended items in a more efficient way. The authors also show that there is an increase of user engagement with the system. User studies of MrTaggy show a similar better user engagement: users spent more time working with the system compared to a baseline system, and had better task results and better understanding of an unfamiliar domain. The authors of CoFeel and Empatheticons evaluated user engagement with subjective feedback from questionnaires and interviews. Results indicate that both systems enhance awareness and user engagement.

Satisfaction and trust have been evaluated by six systems through questionnaires. In general, evaluations focusing on these aspects do indicate some improvement, although the user sample is not always big enough to show the significance. Evaluation results of Bakalov et al. (2013) and PARIS reveal a bias between user control and trustworthiness: the authors indicate that full control over the user profile is not sufficient to establish a good level of trust between the user and the system, as users have privacy concerns. Enabling users to control which data the system can use and for which purposes may reduce these privacy concerns. Results of user studies of TIGRS indicate that trust is strongly related to other evaluation metrics, such as accuracy.

User studies of 15 systems assessed *usefulness* by questionnaires or interviews. Eleven systems evaluated *usability* by post-questionnaires (8), the think aloud method (2) or observing user behavior (1). User studies reveal that visual explanation is useful to help users understand how they get the recommendations (O'Donovan et al., 2008; Bostandjiev et al., 2013; Zhao et al., 2010). In addition, the systems help users to learn more about the underlying data such as similar friends (Gretarsson et al., 2010; Bostandjiev et al., 2012) and relations of recommendations (Parra et al., 2014; Verbert et al., 2013). Bakalov et al. (2013) argues that the integration of the visualization into a recommender system can improve not only its attractiveness, but also the perceived usability.

Whereas user feedback is positive in all cases, some usability issues have been identified. Results of the user studies with TalkExplorer indicate that a cluster map is difficult to use by non-technical users. A comparison with the Venn diagram approach of SetFusion indicated that users are more likely to ex-

plore relationships that help them find useful recommendations (Verbert et al., 2014). A similar comparison of the usefulness of different visualization techniques has been performed by the authors of SmallWorlds. Results indicate that the tree layout works better than the concentric layout, as the layer boundaries are more clear. We elaborate on research opportunities to address these issues in the next section.

6. Challenges

6.1. Objectives

Controllability and transparency have been researched in quite a few systems. But other objectives are still under-explored, including cold start problems, incorporation of contextual information and diversity of recommendations. Some other objectives such as novelty (Herlocker et al., 2004) and serendipity (Herlocker et al., 2004), to the best of our knowledge, have not been tackled explicitly yet with visualizations.

Similar to increasing diversity of recommendations, novelty and serendipity of recommendations focus on suggesting “non-obvious” recommendations. A serendipitous recommendation helps the user find a surprisingly interesting item that she may not have discovered otherwise. Recommendations that are serendipitous are by definition also novel (Herlocker et al., 2004). These factors affect user satisfaction (Tintarev and Masthoff, 2011; Pu et al., 2012; Konstan and Riedl, 2012) and can play an important role in improving current recommender systems.

Although the factors have been studied to some extent by proposing extensions of collaborative filtering techniques (Sarwar et al., 2001), the combination of visualization and recommendation techniques as presented in this article can play a key role to deliver these “non-obvious” recommendations to the user and to support exploration and discovery. An interesting approach has been presented at CHI 2012 (Thudt et al., 2012) that uses visualization to increase diversity of search results. The authors use visualization to offer pathways through digital book collections by providing multiple interactive overviews as visual guides through the collection and by offering many possible adjacencies that can act as visual signposts suggesting alternative exploration routes. Similar support for variety of visual pathways and their flexibility that can serve to enhance serendipity, novelty and diversity of recommendations and constitutes an interesting further line of research.

6.2. Controllability

Controllability of recommendations has been researched extensively over the past decade. Several interesting systems have been surveyed in this article that enable the user to intervene in the recommendation process. Such intervention enables end users to provide input and feedback and is crucial to support the development of a next generation of recommender systems that can be steered by end users. Such a mixed-initiative approach is also promising to address other issues of recommender systems, such as their deployment in high-risk application domains like health-care and financing (McSherry, 2005).

An interesting further line of research is adapting support for user control to different user needs. Previous research shows that the relation between satisfaction and user control is affected by the knowledge level of the user (Knijnenburg et al., 2011a) and her interests (Hijikata et al., 2012). Current interfaces to support user control are *static* - i.e. they do not tailor the interface to these user characteristics. This puts forwards a new topic on interactive recommender systems that can be adapted to different user characteristics and that can support various levels of control in a flexible way. To support such adaptivity, models and techniques that have been used extensively in the adaptive hypermedia research area can be applied (Frias-Martinez et al., 2006). Integrating such approaches with current interactive recommender systems to support adaptive visualization support is promising to advance the current state of the art.

6.3. Context-aware recommendation

Contextual information can be acquired in a number of ways, including explicitly from the user or automatically with sensors. The systems surveyed in this paper use an explicit way to elicit contextual information, and more specifically the current emotion of the user. Visualization techniques are used to enable intuitive acquisition of such variables and to support awareness of such emotional variables in group recommender systems (Chen and Pu, 2013; Chen et al., 2014).

As emotions play a crucial role in decision making (Picard et al., 2004), elaborating this research is of particular interest. In recent years, advancements have been made to acquire information about user emotions in an automatic way with wearable sensors. Our analysis (Reinenbergh et al., 2015) indicates that physiological signals can be used in a successful way to detect different emotions - including happy, sad, fear, disgust and surprise emotions. Khezri et al. (2015) used blood volume pulse, heart rate and skin conductance to measure these emotions. Kim et al. (2005) introduce the measurement of blood pressure, blood volume pulse, skin conductance and skin temperature to detect happy, sad, relaxed and surprise emotions. Whereas both works are very interesting, most of the studies so far are conducted in lab settings (Ouwelkerk, 2011). In our ongoing work, we are using the empatica E4 wristband (Garbarino et al., 2014) to work with these variables in ambient settings. The empatica includes in addition to sensors for detecting blood volume pulse, heart rate variability, skin conductance and skin temperature an accelerometer that can be used to detect whether the user is moving. The approach is promising to start testing whether emotions can be measured in ambient settings, but will no doubt still require input from the user.

Thus, a mixed-initiative approach that enables the user to revise automatically acquired contextual information can again be an interesting future research direction. We are currently researching a combination of that relies on the empatica for automatic detection of emotions and visualization techniques that support awareness and control by end users, for instance to revise detected variables. The overall objective is to research the development of a next generation of recommender systems that can incorporate emotions into the recommendation process.

6.4. Privacy

From a privacy perspective, it is better to let users control whether or not to disclose some piece of information to certain applications and for what purpose (Knijnenburg et al., 2010; Bakalov et al., 2013). Bakalov et al. (2013) shows that privacy concerns are correlated to trust in the system. Explanation interfaces can be an effective method to increase user trust in the system and thereby willingness to disclose personal information (Pu et al., 2012). Research also shows that initial privacy concerns can be overcome when users perceive an improvement of their experience after providing feedback (Knijnenburg et al., 2010). Moreover, privacy concerns reduce once users are highly involved in the system (Spiekermann et al., 2001). Thus, research on user control and explanation interfaces could focus on creating a positive feedback loop that engages users as a basis to increase user trust. In addition to this positive feedback loop, enabling users to control which data can be taken into account for which purposes is a promising research direction for interactive recommender systems, as elaborated in Section 6.2.

6.5. Visualization techniques

Existing interactive recommender systems that we surveyed use node-link diagrams, set-based visualizations, radial views, tables, scatter plots, flow charts and icons. Although the techniques have been shown to work well, there is a need to assess which techniques work better in which settings. Results of our own studies with TalkExplorer indicate that the set-based cluster map technique is too difficult for a non-technical audience. A comparison with a traditional Venn diagram indicates that such a technique is much more suitable for a general audience (Verbert et al., 2014). A similar comparison of the usefulness of different visualization techniques has been performed by the authors of SmallWorlds (Gretarsson et al., 2010), indicating that a tree layout works better than the concentric layout. More generally, there is a need to evaluate which techniques work best under which conditions. Node-link diagrams may for instance work well when the data set is not very large, but they often suffer from visual clutter with there are too many links. Icons may be misleading sometimes (Bogdanov et al., 2013). The rich body of research presented in this paper may serve as a starting point to research design principles and guide researchers in the selection of visualization and interaction techniques. Evaluation of these techniques in different applications and with different end users is key to elaborate design guidelines for a wide audience.

In addition, similar to support for adapting user control as presented in Section 6.2, an interesting future line of research may be adapting visualizations to different user characteristics. More advanced users may benefit from a more complex visualization that is more powerful to gain insight into and interact with recommendation processes. end users with no knowledge of recommender systems and visualization techniques may prefer simple and potentially less advanced visualizations. Providing support for adapting visualizations to the knowledge level and interest of the user is an interesting next step for research in this area.

6.6. Interaction techniques

The work surveyed in this article refers to visualizations developed for desktop or laptop computers, where traditionally the mouse and keyboard are the way the user interacts with the interface. There is an increasing usage of smartphones, tablets, tabletops and surfaces: in 2014, 64% of Americans owned smartphones and 42% a tablet (Pew Research Center, 2014). These multitouch devices have special types of interactions such as tap, drag and pinch that are not available in most desktop or laptop devices. An interesting future line of research is adapting recommender interfaces to different devices, display and interactive technologies.

The works of Song et al. (2013) and Han et al. (2014) show that user interactions available in multitouch devices can be leveraged to increase the accuracy of search and information filtering. Such research opens new opportunities for tailoring interactive visualizations of recommender systems and flexible interactions for use on these devices.

6.7. Evaluation methodology

Finally, there is a need to work with a common evaluation framework that can be used to compare evaluation results of the different systems. Results of our survey indicate that different evaluation methods have been used to assess the impact of visualizations on the recommendation process, ranging from comparisons with baseline data to questionnaires that collect subjective feedback about perceived usefulness and usability. Particularly the latter approach needs to be standardized in order to gain insight into the relative benefits of different approaches and their drawbacks. Such frameworks do exist. A first framework was presented by McNee et al. (2006). Pu et al. (2011) have also presented a promising general user-centric evaluation framework that aims at measuring the quality of recommended items, the usability, usefulness, interface and interaction quality of the system, satisfaction with the systems, and the influence of these qualities on user intentions. Knijnenburg et al. (2011b) have presented a framework to assess subjective system aspects, user experience, interaction and situational and personal characteristics. Uptake of such frameworks would enable to compare among the different techniques and systems and is vital for this research field. The approach would enable to compare and contrast the different approaches that have been presented in this paper and provide a basis for general design guidelines as presented in Section 6.5.

7. Conclusion

In this paper, we have presented an interactive visualization framework of recommender systems that combines recommendation with visualization techniques to enable end users to gain insight into the recommendation process and to help them steer this process. In addition, we have presented an analysis of 24 existing interactive recommender systems along the dimensions of our framework.

Results of this analysis indicate that most existing work focuses on transparency and controllability of the recommendation process. By using visualization techniques, user understanding of the rationale of recommender systems can be supported. Results indicate that such insight can improve acceptance of recommendations. Also supporting user control has an impact on the accuracy of recommendations. Quite a few approaches have been elaborated and shown to perform better compared to a baseline system. Our survey of these approaches in this paper collects many interesting visualization ideas and can guide researchers and practitioners in selecting suitable visualization techniques to support transparency and user control. In addition, the different evaluation approaches that we analyzed may guide these researchers and practitioners to design good user studies and to assess how well these techniques work in a different context.

Although many interesting systems have been elaborated in this research field, there are still many challenges that need to be tackled. First, research objectives such as alleviation of the cold-start problem, diversity, novelty and serendipity of recommendations are still under-explored. Second, there is a need to adapt the level of control and the visualization technique that is used to different user characteristics, as advanced visualizations may be too complex for a wide audience. Third, there is a need to compare and contrast the different techniques with a common evaluation framework and to elaborate design guidelines. The set of visualizations that we have analyzed in this paper presents an interesting starting point with many ideas, but elaborate design guidelines that indicate which techniques are suitable for different users and settings would be helpful to guide new researchers and practitioners. We hope that these ideas can help to further shape exciting and relevant research on interactive recommender systems.

Acknowledgements

Katrien Verbert is a postdoctoral fellow of the Research Foundation Flanders (FWO). Denis Parra was supported by FONDECYT Grant 11150783.

References

- Abdul-Rahman, A., Hailes, S., 2000. Supporting trust in virtual communities. In: Proceedings of the 33rd Annual Hawaii International Conference on System Sciences, p. 9.
- Adnan, W., Daud, N., Noor, N., 2008. Expressive Information Visualization Taxonomy for Decision Support Environment. In: ICCIT08. Third International Conference on Convergence and Hybrid Information Technology. pp. 88–93.
- Adomavicius, G., Tuzhilin, A., 2005. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* 17 (6), 734–749. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1423975>
- Adomavicius, G., Tuzhilin, A., 2012. Context-Aware Recommender Systems. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P. B. (Eds.), *Recommender Systems Handbook*. Springer US, pp. 217–253. URL http://dx.doi.org/10.1007/978-0-387-85820-3_7

- Ahn, J.-W., Brusilovsky, P., 2009. Adaptive visualization of search results: Bringing user models to visual analytics. *Information Visualization* 8 (3), 167–179.
URL <http://www.palgrave-journals.com/doi/10.1057/ivs.2009.12>
- Badenes, H., Bengualid, M. N., Chen, J., Gou, L., Haber, E., Mahmud, J., Nichols, J. W., Pal, A., Schoudt, J., Smith, B. A., et al., 2014. System u: automatically deriving personality traits from social media for people recommendation. In: *Proceedings of the 8th ACM Conference on Recommender systems*. ACM, pp. 373–374.
- Bakalov, F., Meurs, M.-J., König-Ries, B., Sateli, B., Witte, R., Butler, G., Tsang, A., 2013. An approach to controlling user models and personalization effects in recommender systems. In: *Proceedings of the 2013 international conference on Intelligent user interfaces*. ACM, pp. 49–56.
- Bennett, J., Lanning, S., 2007. The netflix prize. In: *Proceedings of KDD cup and workshop*. Vol. 2007. p. 35.
- Bogdanov, D., Haro, M., Fuhrmann, F., Xambó, A., Gómez, E., Herrera, P., 2013. Semantic audio content-based music recommendation and visualization based on user preference examples. *Information Processing & Management* 49 (1), 13–33.
- Bostandjiev, S., O'Donovan, J., Höllerer, T., 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, pp. 35–42.
URL <http://doi.acm.org/10.1145/2365952.2365964>
- Bostandjiev, S., O'Donovan, J., Höllerer, T., 2013. Linkedvis: Exploring social and semantic career recommendations. In: *Proceedings of the 2013 International Conference on Intelligent User Interfaces*. IUI '13. ACM, New York, NY, USA, pp. 107–116.
URL <http://doi.acm.org/10.1145/2449396.2449412>
- Bruns, S., Valdez, A. C., Greven, C., Ziefle, M., Schroeder, U., 2015. What should i read next? a personalized visual publication recommender system. In: *Human Interface and the Management of Information*. Information and Knowledge in Context. Springer, pp. 89–100.
- Burke, R., 2002. Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User Adapted Interaction* 12 (4), 331–370.
URL <http://www.springerlink.com/index/N881136032U8K111.pdf>
- Burke, R., 2010. Evaluating the dynamic properties of recommendation algorithms. In: *Proceedings of RecSys10: fourth ACM conference on Recommender systems*. ACM, New York, NY, USA, pp. 225–228.
URL <http://doi.acm.org/10.1145/1864708.1864753>
- Card, S. K., Mackinlay, J. D., Shneiderman, B., 1999. Using vision to think. Morgan Kaufmann.
- Chen, Y., Ma, X., Cerezo, A., Pu, P., 2014. Empathicons: Designing emotion awareness tools for group recommenders. In: *Proceedings of the XV International Conference on Human Computer Interaction*. ACM, p. 16.
- Chen, Y., Pu, P., 2013. Cofeel: Using emotions to enhance social interaction in group recommender systems. In: *Alpine Rendez-Vous (ARV) 2013 Workshop on Tools and Technology for Emotion-Awareness in Computer Mediated Collaboration and Learning*.
- Chen, Y., Pu, P., 2014. Designing emotion awareness interface for group recommender systems. In: *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces*. ACM, pp. 347–348.
- Chi, E. H., 2000. A taxonomy of visualization techniques using the data state reference model. In: *InfoVis00: Proceedings of the IEEE Symposium on Information Visualization*. pp. 69–75.
- Chuah, M. C., 1998. Dynamic aggregation with circular visual designs. In: *Information Visualization, 1998. Proceedings. IEEE Symposium on*. IEEE, pp. 35–43.
- Dey, A. K., 2001. Understanding and using context. *Personal and ubiquitous computing* 5 (1), 4–7.
- Dourish, P., Feb. 2004. What we talk about when we talk about context. *Personal Ubiquitous Computing* 8 (1), 19–30.
URL <http://dx.doi.org/10.1007/s00779-003-0253-8>
- Ellis, G., Dix, A., Nov. 2007. A Taxonomy of Clutter Reduction for Information Visualisation. *IEEE Transactions on Visualization and Computer Graphics* 13 (6), 1216–1223.
URL <http://dx.doi.org/10.1109/TVCG.2007.70535>
- Elmqvist, N., Fekete, J., 2010. Hierarchical Aggregation for Information Visualization: Overview, Techniques, and Design Guidelines. *IEEE Transactions on Visualization and Computer Graphics* 16 (3), 439–454.
- Elmqvist, N., Stasko, J., Tsigas, P., 2008. DataMeadow: A Visual Canvas for Analysis of Large-Scale Multivariate Data. *Information Visualization* 7 (1), 18–33.
- Fekete, J.-D., Van Wijk, J. J., Stasko, J. T., North, C., 2008. The Value of Information Visualization. *Information Visualization*, 1–18.
- Felfernig, A., Gula, B., 2006. An empirical study on consumer behavior in the interaction with knowledge-based recommender applications. In: *E-Commerce Technology, 2006. The 8th IEEE International Conference on and Enterprise Computing, E-Commerce, and E-Services, The 3rd IEEE International Conference on*. IEEE, pp. 37–37.
- Frias-Martinez, E., Chen, S. Y., Liu, X., 2006. Survey of data mining approaches to user modeling for adaptive hypermedia. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* 36 (6), 734–749.
- Garbarino, M., Lai, M., Bender, D., Picard, R. W., Tognetti, S., 2014. Empatica e3a wearable wireless multi-sensor device for real-time computerized biofeedback and data acquisition. In: *Wireless Mobile Communication and Healthcare (MobiHealth), 2014 EAI 4th International Conference on*. IEEE, pp. 39–42.
- Gou, L., You, F., Guo, J., Wu, L., Zhang, X. L., 2011. Sfviz: interest-based friends exploration and recommendation in social networks. In: *Proceedings of the 2011 Visual Information Communication-International Symposium*. ACM, p. 15.
- Gretarsson, B., O'Donovan, J., Bostandjiev, S., Hall, C., Höllerer, T., 2010. SmallWorlds: Visualizing social recommendations. *Computer Graphics Forum* 29 (3), 833–842.
- Halder, S., Samiullah, M., Lee, Y.-K., et al., 2012. Movie swarm: Information mining technique for movie recommendation system. In: *Electrical & Computer Engineering (ICEE), 2012 7th International Conference on*. IEEE, pp. 462–465.
- Han, S., Hsiao, I.-H., Parra, D., 2014. A study of mobile information exploration with multi-touch interactions. In: *Social Computing, Behavioral-Cultural Modeling and Prediction*. Springer, pp. 269–276.
- Herlocker, J. L., Konstan, J. A., Riedl, J., 2000. Explaining collaborative filtering recommendations. *Proceedings of CSCW00: ACM conference on Computer supported cooperative work pages*, 241–250.
URL <http://portal.acm.org/citation.cfm?doid=358916.358995>
- Herlocker, J. L., Konstan, J. A., Terveen, L. G., Riedl, J. T., 2004. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems* 22 (1), 5–53.
URL <http://portal.acm.org/citation.cfm?doid=963770.963772>
- Herman, I., Melancon, G., Marshall, M. S., 2000. Graph visualization and navigation in information visualization: A survey. *Visualization and Computer Graphics, IEEE Transactions on* 6 (1), 24–43.
- Hijikata, Y., Kai, Y., Nishida, S., 2012. The relation between user intervention and user satisfaction for information recommendation. In: *Proceedings of the 27th Annual ACM Symposium on Applied Computing*. SAC '12. ACM, New York, NY, USA, pp. 2002–2007.
URL <http://doi.acm.org/10.1145/2245276.2232109>
- Hong, J., Suh, E.-H., Kim, J., Kim, S., 2009. Context-aware system for proactive personalized service based on context history. *Expert Systems with Applications* 36 (4), 7448–7457.
- Hu, R., Pu, P., 2011. Helping users perceive recommendation diversity. In: *Workshop on novelty and diversity in recommender systems, DiveRS*. Chicago.
- Inselberg, A., Dec. 1985. The plane with parallel coordinates. *The Visual Computer* 1 (4), 69–91.
- Jin, y., Karsten, S., Verbert, K., Duval, E., 2015. Effects of transparency and user control on targeted advertising (submitted to IUI 2016). <http://paris-ad.evennode.com/paris/>.
- Kammerer, Y., Nairn, R., Piroli, P., Chi, E. H., 2009. Signpost from the masses: learning effects in an exploratory social tag search browser. In: *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, pp. 625–634.
- Kangasrääsiö, A., Glowacka, D., Kaski, S., 2015. Improving controllability and predictability of interactive recommendation interfaces for exploratory search. In: *Proceedings of the 20th International Conference on Intelligent User Interfaces*. ACM, pp. 247–251.

- Keim, D. A., 2002. Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics* 8 (1), 1–8.
URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=981847>
- Khezri, M., Firoozabadi, M., Sharafat, A. R., 2015. Reliable emotion recognition system based on dynamic adaptive fusion of forehead biopotentials and physiological signals. *Computer Methods and Programs in Biomedicine*.
- Kim, S.-J., Kim, J.-S., Kim, S.-H., Kim, Y.-M., 2005. Evolvable recommendation system in the portable device based on the emotion awareness. In: *Knowledge-Based Intelligent Information and Engineering Systems*. Springer, pp. 251–257.
- Klerkx, J., Duval, E., 2009. Visualising social bookmarks. *Journal of Digital Information* 10 (2), 1–40.
URL <https://lirias.kuleuven.be/handle/123456789/234788>
- Knijnenburg, B. P., Reijmer, N. J., Willemsen, M. C., 2011a. Each to his own: How different users call for different interaction methods in recommender systems. In: *Proceedings of the Fifth ACM Conference on Recommender Systems*. RecSys '11. ACM, New York, NY, USA, pp. 141–148.
URL <http://doi.acm.org/10.1145/2043932.2043960>
- Knijnenburg, B. P., Willemsen, M. C., Hirtbach, S., 2010. Receiving recommendations and providing feedback: The user-experience of a recommender system. In: *E-Commerce and Web Technologies*. Springer, pp. 207–216.
- Knijnenburg, B. P., Willemsen, M. C., Kobsa, A., 2011b. A pragmatic procedure to support the user-centric evaluation of recommender systems. In: *Proceedings of the fifth ACM conference on Recommender systems*. ACM, pp. 321–324.
- Konstan, J., Riedl, J., 2012. Recommender systems: from algorithms to user experience. *User Modeling and User-Adapted Interaction* 22 (1-2), 101–123.
URL <http://dx.doi.org/10.1007/s11257-011-9112-x>
- Lamping, J., Rao, R., 1996. The Hyperbolic Browser: A Focus+Context Technique for Visualizing Large Hierarchies. *Journal of Visual Languages & Computing* 7 (1), 33–55.
URL <http://www.sciencedirect.com/science/article/B6WMM-45PVN1K-1N/2/df3d57633ecd6d45f41986fc05c8fbd8>
- Lerner, J. S., Li, Y., Valdesolo, P., Kassam, K. S., 2015. Emotion and decision making. *Psychology* 66.
- Loepp, B., Hussein, T., Ziegler, J., 2014. Choice-based preference elicitation for collaborative filtering recommender systems. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, pp. 3085–3094.
- McNee, S. M., Riedl, J., Konstan, J. A., 2006. Making recommendations better: an analytic model for human-recommender interaction. In: *CHI '06 extended abstracts on Human factors in computing systems*. ACM, pp. 1103–1108.
- McSherry, D., 2005. Explanation in recommender systems. *Artificial Intelligence Review* 24 (2), 179–197.
- Morville, P., 2005. *Ambient Findability*. O'Reilly Media, Inc.
- Nageswara, R. K., Talwar, V. G., 2008. Application domain and functional classification of recommender systems: a survey. *DESIDOC Journal of Library and Information Technology* 28 (3), 17–35.
- Norman, D. A., 2002. The psychopathology of everyday things. *Foundations of Cognitive Psychology: Core Readings*, 417.
- Nowell, L., Havre, S., Hetzler, E., Whitney, P., 2002. Themeriver: Visualizing Thematic Changes in Large Document Collections. *IEEE Transactions on Visualization and Computer Graphics* 8 (1), 9–20.
- Nussbaumer, A., 2005. Integrating Four Graph-Based Hierarchy Browsers into the Hierarchical Visualisation System (HVS). Master thesis, Graz University of Technology.
- O'Donovan, J., Smyth, B., Gretarsson, B., Bostandjiev, S., Höllerer, T., 2008. PeerChooser: Visual Interactive Recommendation. In: *Proceedings of CHI08: twenty-sixth annual SIGCHI conference on Human factors in computing systems*. ACM, pp. 1085–1088.
URL <http://portal.acm.org/citation.cfm?id=1357222>
- Ouwerkerk, M., 2011. Unobtrusive emotions sensing in daily life. In: *Sensing Emotions*. Springer, pp. 21–39.
- Parra, D., Brusilovsky, P., 2015. User-controllable personalization: A case study with setfusion. *International Journal of Human Computer Studies* 78, 43–67.
- Parra, D., Brusilovsky, P., Trattner, C., 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, pp. 235–240.
URL <http://doi.acm.org/10.1145/2557500.2557542>
- Pazzani, M. J., Billsus, D., 2007. Content-Based Recommendation Systems. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (Eds.), *The Adaptive Web*. Vol. 4321 of LNCS. Springer-Verlag, Springer, pp. 325–341.
URL <http://www.springerlink.com/index/qq35wt6816774261.pdf>
- Pew Research Center, 2014. Mobile technology fact sheet. <http://www.pewinternet.org/fact-sheets/mobile-technology-fact-sheet/>.
- Picard, R. W., Papert, S., Bender, W., Blumberg, B., Breazeal, C., Cavallo, D., Machover, T., Resnick, M., Roy, D., Strohecker, C., 2004. Affective learning - a manifesto. *BT Technology Journal* 22 (4), 253–269.
- Pu, P., Chen, L., Hu, R., 2011. A user-centric evaluation framework for recommender systems. In: *Proceedings of the fifth ACM conference on Recommender systems*. ACM, pp. 157–164.
- Pu, P., Chen, L., Hu, R., 2012. Evaluating recommender systems from the users perspective: survey of the state of the art. *User Modeling and User-Adapted Interaction* 22 (4-5), 317–355.
- Reinenbergh, J., Karsten, S., Verbert, K., 2015. Emotion-aware recommender systems. <https://jeroenrthesis.wordpress.com/>.
- Saito, Y., Itoh, T., 2011. Musicube: a visual music recommendation system featuring interactive evolutionary computing. In: *Proceedings of the 2011 Visual Information Communication-International Symposium*. ACM, p. 5.
- Santos, R. D., Gros, P., Abel, P., Loisel, D., Trichaud, N., Paris, J., 2000. Experiments in Information Visualization Using 3D Metaphoric Worlds. In: *Proceedings of WETICE00: Enabling Technologies: Infrastructure for Collaborative Enterprises*. pp. 51–58.
URL http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=863704
- Sarwar, B., Karypis, G., Konstan, J., Riedl, J., 2001. Item-based collaborative filtering recommendation algorithms. In: *Proceedings of the 10th international conference on World Wide Web*. ACM, pp. 285–295.
- Schaffner, J., Höllerer, T., O'Donovan, J., 2015. Hypothetical recommendation: A study of interactive profile manipulation behavior for recommender systems. In: *The Twenty-Eighth International Flairs Conference*.
- Schein, A. I., Popescul, A., Ungar, L. H., Pennock, D. M., 2002. Methods and metrics for cold-start recommendations. In: *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, pp. 253–260.
- Scherer, K. R., 2005. What are emotions? and how can they be measured? *Social science information* 44 (4), 695–729.
- Shneiderman, B., 1996. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In: *Proceedings of VL96: IEEE Symposium on Visual Languages*. IEEE Computer Society, pp. 336–343.
- Shneiderman, B., Bederson, B. B., 2003. *The Craft of Information Visualization: Readings and Reflections*. Morgan Kaufmann.
- Shneiderman, B., Johnson, B., 1991. Tree-Maps: A Space Filling Approach to the Visualization of Hierarchical Information Structures. In: *Proceedings of VIS91: Second Conference on Visualization*. pp. 284–291.
- Sinha, R., Swearingen, K., 2002. The role of transparency in recommender systems. In: *CHI '02 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '02. ACM, New York, NY, USA, pp. 830–831.
URL <http://doi.acm.org/10.1145/506443.506619>
- Song, Y., Ma, H., Wang, H., Wang, K., 2013. Exploring and exploiting user search behavior on mobile and tablet devices to improve search relevance. In: *Proceedings of the 22nd international conference on World Wide Web*. International World Wide Web Conferences Steering Committee, pp. 1201–1212.
- Spiekermann, S., Grossklags, J., Berendt, B., 2001. E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In: *Proceedings of the 3rd ACM conference on Electronic Commerce*. ACM, pp. 38–47.
- Stasko, J., Zhang, E., 2000. Focus+ context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In: *Information Visualization, 2000. InfoVis 2000. IEEE Symposium on*. IEEE, pp. 57–65.
- Stefaner, M., Urban, T., Marc, S., 2008. Elastic Lists for Facet Browsing and Resource Analysis in the Enterprise. In: *Proceedings of DEXA08: Nineteenth International Conference on Database and Expert Systems Applications*. IEEE Computer Society, pp. 397–401.
- Swearingen, K., Sinha, R., 2001. Beyond Algorithms: An HCI Perspective on Recommender Systems. *ACM SIGIR 2001 Workshop on Recommender Systems*, 1–11.

- Symeonidis, P., Nanopoulos, A., Manolopoulos, Y., 2009. MoviExplain: A Recommender System with Explanations. Proceedings of the third ACM conference on Recommender systems - RecSys '09, 317–320.
URL <http://portal.acm.org/citation.cfm?doid=1639714.1639777>
- Ternier, S., Verbert, K., Parra, G., Vandeputte, B., Klerkx, J., Duval, E., Ordonez, V., Ochoa, X., Jul. 2009. The Ariadne Infrastructure for Managing and Storing Metadata. IEEE Internet Computing 13 (4), 18–25.
URL <http://www.computer.org/portal/web/csd1/doi/10.1109/MIC.2009.90>
- Thudt, A., Hinrichs, U., Carpendale, S., 2012. The bohemian bookshelf: supporting serendipitous book discoveries through information visualization. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, pp. 1461–1470.
- Tintarev, N., Masthoff, J., 2011. Designing and evaluating explanations for recommender systems. In: Recommender Systems Handbook. Springer, pp. 479–510.
- Van Wijk, J. J., van Seelow, E. R., 1999. Cluster and Calendar Based Visualization of Time Series Data. In: Proceedings of the IEEE Symposium on Information Visualization. pp. 4–9.
- Verbert, K., Manouselis, N., Ochoa, X., Wolpers, M., Drachsler, H., Bosnic, I., Duval, E., 2012. Context-aware recommender systems for learning: a survey and future challenges. Learning Technologies, IEEE Transactions on 5 (4), 318–335.
- Verbert, K., Parra, D., Brusilovsky, P., 2014. The effect of different set-based visualizations on user exploration of recommendations. In: Proceedings of the Joint Workshop on Interfaces and Human Decision Making in Recommender Systems. pp. 37–44.
- Verbert, K., Parra, D., Brusilovsky, P., Duval, E., 2013. Visualizing recommendations to support exploration, transparency and controllability. Proceedings of the 2013 international conference on Intelligent user interfaces - IUI '13, 351.
URL <http://dl.acm.org/citation.cfm?id=2449396.2449442>
- Vig, J., Riedl, J., 2009. Tagsplanations : Explaining Recommendations Using Tags. Work, 47–56.
URL <http://portal.acm.org/citation.cfm?id=1502650.1502661>
- Vlachos, M., Svonava, D., 2012. Graph embeddings for movie visualization and recommendation. CEUR Workshop Proceedings 891, 41–47.
- Ware, C., 2000. Information visualization: perception for design. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Willett, W., Heer, J., Agrawala, M., 2007. Scented Widgets: Improving Navigation Cues with Embedded Visualizations. IEEE Transactions on Visualization and Computer Graphics 13 (6), 1129–1136.
URL <http://dblp.uni-trier.de/db/journals/tvcg/tvcg13.html#WillettHA07>
- Wong, D., Faridani, S., Bitton, E., Hartmann, B., Goldberg, K., 2011. The diversity donut: enabling participant control over the diversity of recommended responses. In: CHI'11 Extended Abstracts on Human Factors in Computing Systems. ACM, pp. 1471–1476.
- Yau, J., Joy, M., 2007. A context-aware and adaptive learning schedule framework for supporting learners' daily routines. In: Systems, 2007. ICONS'07. Second International Conference on. IEEE, pp. 31–31.
- Yi, J. S., Elmqvist, N., Lee, S., 2010. TimeMatrix: Analyzing Temporal Social Networks Using Interactive Matrix-Based Visualizations. International Journal of Human-Computer Interaction 26 (11&12), 1031–1051.
- Zhao, S., Zhou, M. X., Yuan, Q., Zhang, X., Zheng, W., Fu, R., 2010. Who is Talking about What : Social Map-based Recommendation for Content-Centric Social Websites. 4th ACM Conference on Recommender Systems, 143–150.
- Zheng, Y., Mobasher, B., Burke, R. D., 2013. The role of emotions in context-aware recommendation. Decisions@ RecSys 2013, 21–28.