Reading News with a Purpose: Explaining User Profiles for Self-Actualization

Emily Sullivan*, Dimitrios Bountouridis*, Jaron Harambam[†], Shabnam Najafian*, Felicia Loecherbach^{††},
Mykola Makhortykh[‡], Domokos Kelen^{‡‡}, Daricia Wilkinson^{||}, David Graus^{**}, Nava Tintarev*

*Delft University of Technology, The Netherlands

[†]Institute for Information Law, University of Amsterdam, The Netherlands

[‡]Amsterdam School of Communication Research, University of Amsterdam, The Netherlands

^{††}Vrije University, Amsterdam, The Netherlands

^{‡‡} Institute for Computer Science and Control, Hungarian Academy of Sciences, Hungary

^{||}Clemson University, Clemson, USA

^{**}FD Mediagroep

e.e.sullivan-mumm@tudelft.nl

ABSTRACT

Personalized content provided by recommender systems is an integral part of the current online news reading experience. However, news recommender systems are criticized for their 'black-box' approach to data collection and processing, and for their lack of explainability and transparency. This paper focuses on explaining user profiles constructed from aggregated reading behavior data, used to provide content-based recommendations. The paper makes a first step toward consolidating epistemic values of news providers and news readers. We present an evaluation of an explanation interface reflecting these values, and find that providing users with different goals for self-actualization (i.e., *Broaden Horizons* vs. *Discover the Unexplored*) influences their reading intentions for news recommendations.

KEYWORDS

explainability, user profile, user control, self-actualization, news recommender systems

ACM Reference Format:

Emily Sullivan, Dimitrios Bountouridis, Jaron Harambam, Shabnam Najafian, Felicia Loecherbach, Mykola Makhortykh, Domokos Kelen, Daricia Wilkinson, David Graus, & Nava Tintarev. 2019. Reading News with a Purpose: Explaining User Profiles for Self Actualization. In *Proc. of 27th Conf. on User Modeling, Adaption & Personalization Adjunct (UMAP'19 Adjunct)*, June 9–12, 2019, Larnaca, Cyprus. ACM, NY, NY. 5 pages.
<https://doi.org/10.1145/3314183.3323456>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UMAP'19 Adjunct, June 9–12, 2019, Larnaca, Cyprus

© 2019 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-6711-0/19/06...\$15.00

<https://doi.org/10.1145/3314183.3323456>

1 INTRODUCTION

Personalized experiences powered by recommender systems permeate our lives. However, the rise of distrust and skepticism around the collection and use of personal data generates an increased interest in *explainability* and *transparency* of black-box recommender systems [12, 13, 31]. Many works [14, 19, 35] show that explanations are integral for users to understand their consumption preferences and achieve their *epistemic goals* i.e., goals for knowledge development. Moreover, researchers have begun exploring methods that support the exploration and understanding of users' current, aspirational, and self-actualized goals to provide transparency in recommender systems [11, 22]. Yet, the task of opening the black box of recommender systems remains notoriously hard to achieve [8, 28]. For the domain of *online news*, this might limit the readers' control over information diets [15]. Limited explainability can also diminish users' engagement with recommended content [21], as well as raising concerns about the societal role of news media [9] and consequences of increased algorithmization [26, 33].

To address these worries, this paper focuses on the goal of explaining recommendations in the context of *online news*. We explore methods of explaining one aspect of how a content-based recommenders work: the *user profile*. We aim to automatically summarize and visualize the recommender's high dimensional internal representations of users. These profiles are automatically generated from their reading behavior, through leveraging news article's topics, entities, and tags. Explaining user profiles have been demonstrated to be effective [3], particularly relevant for the news domain [14] and able to facilitate users' self-actualization [22].

When approaching the problem of explaining user profiles, questions such as "what is the purpose of the explanations", "what user-goal do they serve" or "what type of user-control and visualization should they include" immediately emerge. As such, this paper's contribution is twofold: (i) we organize the concepts of explanations, user control and user goals in a systematic fashion. In Section 3, we present a conceptual *framework* that defines goals of explanations for the news domain. We distinguish between three layers of explanations that enable users to answer different questions depending on their own and the content provider's epistemic goals

for self-actualization. (ii) We instantiate the proposed framework by *evaluating* an explanation interface for two self-actualization goals in Section 4. We study whether *providing users with different goals* influences which topics they intend to read next. We now first turn to a discussion of related work on transparency and user control in recommender systems.

2 RELATED WORK

Our approach to explaining news recommenders is based on individual user profiles. Middleton et al. [24] first visualized user interests as time/interest graphs. These were demonstrated to be effective in a commercial setting by Bonhard and Sasse [3]. Tintarev and Masthoff [36] identify several groups of explanation styles depending on algorithms' types and the domains they are used in. These styles vary depending on explanations' focus (e.g. content- or utility based explanations) or specific format (e.g. conversational or demographic explanations). In the online news domain, the majority of studies focus on content-based mechanisms using automatically generated textual explanations [2, 32, 34].

It has been suggested that user-control (the ability of the user to provide input or feedback affecting the system) can be applied to different parts of the recommendation pipeline [17, 19]. The first part of the pipeline corresponds to preference elicitation i.e., users explicitly specifying their preferences either via static profile forms [19] or more dynamic solutions such as the MusiCube interface [30]. The second part corresponds to choosing or influencing recommendation strategy [10]. For the latter, several control mechanisms exist, from draggable sliders and Venn diagrams [27], to node-link diagrams [6] and Likert scales [20]. Finally, the third part filters the recommender's output typically via mechanics such as list-reranking. A number of works have presented systems that provide control on more than one layer of the recommendation pipeline, such as TasteWeights [4] or LinkedVis [5]. A more complete review can be found in He et al. [17].

3 EXPLAINABILITY FRAMEWORK

User profiles can be used as explanations, and depending on the question asked different explanations may be appropriate. We adopt an approach that is a useful departure from prior methods around the explainability of algorithmic systems in general [1, 41], and in the news domain [9]. We ask: *How can we use these various functions of explanations to increase transparency and promote self-actualization?*

3.1 Levels of explanation

We identified three generalizable functions, *levels of explanation* that a user profile can serve. Each level takes the information from the previous level and restructures it in a way that, we argue, steadily increases transparency and understanding.

Level 1: Transparency. The first level consists of the raw data that the platform has on the user and the user's reading history. This data serves as the first layer since it provides the information necessary for more higher level questions or goals that the user might have. On the one hand, the transparency layer serves simply the function of transparency; however, this data can also help the user answer information seeking questions regarding their past

reading behavior (e.g., by visualizing the distribution of monthly read topics by the user).

Level 2: Contextualization. The second layer adds to the first by taking a user's specific past behavior and contextualizing it within their community. This layer helps the user understand how others are using the news platform (e.g., by visualizing side-by-side the distribution of monthly read topics by the current and the average user). This has been found to help users answer questions about how their consumption habits compare to the overall user base [23, 37]. Other possible comparisons include, comparing a particular user's consumption to those with similar consumption habits or to the platform's publication history. It has been previously noted however that explanations based on a global profile (all users) can also potentially lead to undesirable consequences, such as nudging users to become closer to the mean [37].

Level 3: Self-Actualization. The third layer supports epistemic goals that foster self-actualization. Few studies have tried to connect transparency and explanation with certain personal or societal values and goals [15, 22, 23, 38]. This layer departs from the simple information-finding of previous layers by promoting discovery and exploration. It is *goal-directed* and allows for *user-control* to achieve those goals.

In this layer, the user has direct *control* over which goal they want to explore, and the recommendations that result from the chosen goal. The user should be presented not only with the different goals, but also with a short textual explanation of how exploring the goal can help them achieve their epistemic goals. In our topic-distribution example, after the user chooses a goal, they can explore new recommended topics through an interactive visualization and a user-control mechanism.

3.2 Self-actualization goals

In collaboration with a Dutch online news platform (FD Mediagroep's "Het Financieele Dagblad", fd.nl) we identified two goals aligning with FD Mediagroep's values that support self-actualization:

Broaden Horizons: This goal highlights the need for diversity in news selection [18]. Broadening one's horizon can help combat polarization, in addition to achieving greater personal well-being, understanding, and tolerance [7]. Moreover, in the high-choice information environment, assessing which perspective is actually truthful can be difficult [16]. Achieving a broadening of horizons is possible by showing users articles, perspectives and topics they normally do not read about, and that they may find interesting. We do not connect users to completely unrelated items, since there may be well-founded reasons why they have left them untouched. The idea is to gradually increase diversity, to slowly move out of their epistemic bubble [39].

Discover the Unexplored: Discovering the unexplored utilizes long-tail diversity of the news platform. This goal is inspired by the notion of serendipity [29]. Ordinary news reading behavior is characterized by browsing through the news, and accidentally stumbling upon content that attract one's attention. When news recommenders are attuned to accuracy only, it becomes difficult to come across topics that one has not read or thought of before,

but that may actually be(come) of interest. Discovering the unexplored helps users encounter unknown articles that may spark new interests or provide new contexts. Effectively bringing users to completely new territories, out of their comfort zone. In this vein, a previous study found that visualizing under-explored parts of a user profile improves exploration in the music domain [23].

4 USER STUDY ON SELF-ACTUALIZATION

To validate part of our framework, we focus on Level 3. In this user study we want to explore *whether a user being prompted with a particular goal would influence their intended reading behavior*.

We identified two self-actualization goals: *Broaden Horizons* and *Discover the Unexplored*. Given the relatedness between these two goals—both involve users exploring diverse content—it serves as a suitable test case as to whether the specific goal leads to the user to explore different degrees of diversity. Moreover, even though we expect that the most suitable visualization differs for each goal, we fix the visualization for both, thus ensuring that the users' choices are only impacted by the stated goal. The user study was designed to test the following hypotheses:

H1: Goal Framework. Providing the user with a stated goal will influence which topics they wish to read about next, in *Broaden Horizons* they select more similar (**H1a**) and familiar (**H1b**) topics compared to *Discover the Unexplored*.

H2: Broaden Horizons Familiarity. We expect users to choose topics that have high similarity (**H2a**), and high familiarity (**H2b**), compared to their non-selected topics, since the goal of broadening horizons is to make small steps outside of current reading behavior.

H3: Discover the Unexplored Familiarity. We expect users to choose topics that have low similarity (**H3a**), and low familiarity (**H3b**), compared to their non-selected topics, since the goal of discover the unexplored urges users to explore topics they are largely outside of their current interests.

4.1 Materials

Data. We were provided with a month of real data of reading behavior of users of fd.nl, the online news portal of FD Media-groep's newspaper "Het Financieel Dagblad." The dataset consists of a sample of 100 user profiles, their reading behavior (during November 2018), and the metadata of the articles (over 1600 unique articles) that were read by these users during the sampled month. The 100 users were sampled by reading activity. The users were binned by the number of articles read, and we were given a random sample of 50 users with *average* reading activity (i.e., sampled from the middle bin), and 50 *highly active* users (i.e., sampled from the highest bin). Of these, we extracted four user profiles that showed a maximum distance to each other, as described in Section 4.2. The user profiles were fully anonymized: user IDs were hashed, and identifiable information such as users' names, email addresses, locations, etc. were not part of the data set. For this study we restricted to only using the article's tags (manually added by editors) and the averages of the word embeddings from the article content.

Visualization. We used a visualization that shows the interests of the user based on their past reading behavior in relation to the

Figure 1: User profile. Nodes are news topics, ordered by the user's familiarity with them, and darker color indicates more familiarity. The connections between topics represent their similarity, with thicker lines meaning more similar.

publication history of the platform, and also the similarities between these topics and other unexplored ones (see Figure 1). It allows the user to identify topics they are familiar with; topics that are close to their interests, but are largely unexplored, and topics that are completely missing from their reading history. In a real application, the user would be able to interact with the topics (e.g., clicking them) to influence the recommender system.

The connections between topics were determined by language similarity: we used the cosine similarity of aggregated word embeddings [25] of the articles in the given topics, and kept connections above an arbitrarily-defined threshold.

4.2 Experimental Design

Respondents first answered questions regarding their demographics (age, gender, education).

Profile matching. Respondents then indicated their news topic interests by rating how often (5-point scale from 1 'never' to 5 'always') they read about seven topics (politics, foreign affairs, economy, lifestyle, sports, technology, arts). Afterwards, they were presented with four different word clouds representing user profiles, and chose which one best matched their interests.

We extracted four user profiles that showed a maximum distance to each other (based on the similarity of the tags they used) to ensure broad variation in profiles. The Dutch tags were translated and adapted for American respondents (i.e., removing information specific to the Netherlands such as names of politicians or party and company names). The translated tags read by those users were transformed into word clouds. This process captured the four distinct user profiles respondents could choose from: technology and media, EU politics, US politics, and the economy.

Explanation goals. There were two experimental conditions: a) the broaden horizon condition, and b) the discover the unexplored condition. Each respondent answered questions for both conditions so analysis within and across respondents was possible. The order of conditions was randomized for each respondent to minimize the effect of order bias in our results.

In the *discover the unexplored condition* respondents were given the following goal:

Your goal is to Discover the Unexplored. There may be topics that you haven't explored before that may actually become new interests. Exploring new topics can promote creativity and objectivity.

In the *broaden horizon condition* respondents were given the following goal:

Your goal is to Broaden your Horizons. There may be topics you do not normally read about, but you may actually find interesting. Exploring this helps to build a broad perspective on the issues that matter to you.

After the goal was introduced respondents were shown the visualization corresponding to the user profile they chose and asked to pick from a drop down menu (containing all the topics represented in the visualization) which topics they want to explore next: “Which three topics do you want to explore for [goal]? Please select three.”

4.3 Topic Scoring

We test our hypotheses by computing user familiarity and similarity scores for topics.

Familiarity. The user's familiarity score with a topic is computed as the ratio of articles read by a user on a topic, over the total number of articles published on that topic. We compute the average familiarities of the three topics selected by the user, and compare that to the average familiarity of the non-selected topics using the Wilson score confidence interval [40].

Similarity. The similarity between topics in the user profile is computed using the cosine similarity of aggregated word embeddings [25] of the articles in the given topics. We compute the average similarity of the three topics selected by the user, to the set of topics a user is familiar with (i.e., topics where familiarity score > 0), and compare that to the average similarity of the non-selected topics to the topics the user is familiar with.

4.4 Results

Participants. Fifty eight respondents were recruited from Amazon Mechanical Turk (MTurk), on January 25, 2019. We recruited respondents from the U.S. who had achieved the qualification as ‘master’ (reliable worker) to ensure a high quality of data collection. **Age:** 3% 18-24; 41% 25-34; 34% 35-44; 11% 45-54; 11%; 55 or older. **Gender:** 54% male; 43% female; 3% other. **Education:** 2% less than high school; 16% high school or equivalent; 29% some college, no degree, 41% bachelor degree, 11% graduate degree.

In analyzing each hypothesis, we used the Wilcoxon signed-rank test, and applied Bonferroni correction.

4.4.1 H1. Goal Framework.

- **H1a. Similarity.** In the Broaden Horizons condition users selected more similar topics ($M=54.4$, $SD=25.8$) compared to Discover the Unexplored condition ($M=51.7$, $SD=28.1$), but this trend is not statistically significant ($Z = -0.1$, $p = 0.9$).
- **H1b. Familiarity.** In Broaden Horizons condition users selected more familiar topics ($M=1.01$, $SD=0.02$) compared to Discover the Unexplored condition ($M=1.0$, $SD=0.01$). The result was statistically significant ($Z = -3.6$, $p < 0.01$).

While H1a is not supported, H1b is accepted: *Providing users with a stated goal does influence which topics they wish to read about next; for familiarity.*

4.4.2 H2. Broaden Horizons.

- **H2a. Similarity.** The means for the selected ($M=54.4$, $SD=25.8$) is slightly lower than non-selected topics ($M=57.1$, $SD=2.6$); however the variance of similarity for selected topics is much higher. The difference between average similarity of selected and non-selected topics is not significant ($Z = -0.4$, $p = 1.4$).
- **H2b. Familiarity.** The means for the selected ($M=1.01$, $SD=0.02$) and non-selected topics ($M=1.01$, $SD=0.01$) is similar. The difference between average familiarity of selected and non-selected topics is not significant ($Z = -1.0$, $p = 0.6$).

Both H2a and H2b are rejected. For Broaden Horizons, there is no difference between selected and non-selected topics w.r.t familiarity or similarity.

4.4.3 H3. Discover the Unexplored.

- **H3a. Similarity.** Users selected less similar topics ($M=51.7$, $SD=28.1$) compared to non-selected topics ($M=57.0$, $SD=1.8$), with a larger variance for selected topics. This result was not statistically significant ($Z = -0.8$, $p = 0.4$).
- **H3b. Familiarity.** Users selected less familiar topics ($M=1$, $SD=0.01$) compared to non-selected topics ($M=1.01$, $SD=0.02$). This result was statistically significant ($Z = -5.2$, $p < 0.01$).

We find partial support for H3b. *When presented the goal of Discover the Unexplored, users selected less familiar topics.*

Discussion. Overall, the results of our user study suggest that providing users with a clear goal influences their reading intentions (stopping short of studying actual reading behaviors). This is true especially regarding topic familiarity.

The mixed results within the broaden horizon condition could be due to the fact that the visualization itself was not focused on this specific goal. Moreover, we approximated the user preferences using someone else's profile. Respondents might have chosen topics that cohered with broadening their own interests more so than the selected profile. Respondents might have an easier time selecting far removed topics (the goal of discovering the unexplored) compared to selecting nearby topics for the purpose of broadening horizons.

5 CONCLUSION

Motivated by the lack of explainability and transparency in news recommender systems, this paper focused on explaining the user's news profile. We presented a novel and generalizable three-layered framework of epistemic goals and associated types of explanation. Focusing on a single layer, this paper presented a first encouraging step toward consolidating two specific values of a news provider and its potential readers. In future work we plan to develop visualizations that better complement each specific goal, and study how these goals can be applied to influence actual reading behaviors.

ACKNOWLEDGEMENTS

We would like to thank ICT with Industry, NWO, and FD Media-groep.

REFERENCES

- [1] Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society* 20, 3 (2018), 973–989.
- [2] Roi Blanco, Diego Ceccarelli, Claudio Lucchese, Raffaele Perego, and Fabrizio Silvestri. 2012. You should read this! let me explain you why: explaining news recommendations to users. In *Proceedings of the 21st ACM international conference on Information and knowledge management*. ACM, 1995–1999.
- [3] Philip Bonhard and Martina Angela Sasse. 2006. "Knowing me, knowing you"–Using profiles and social networking to improve recommender systems. *BT Technology Journal* 24, 3 (2006), 84–98.
- [4] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. TasteWeights: a visual interactive hybrid recommender system. In *Proceedings of the sixth ACM conference on Recommender systems*. ACM, 35–42.
- [5] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2013. LinkedVis: exploring social and semantic career recommendations. In *Proceedings of the 2013 international conference on Intelligent user interfaces*. ACM, 107–116.
- [6] Simon Bruns, André Calero Valdez, Christoph Greven, Martina Ziefle, and Ulrik Schroeder. 2015. What should I read next? a personalized visual publication recommender system. In *International Conference on Human Interface and the Management of Information*. Springer, 89–100.
- [7] Nicole Curato, John S Dryzek, Selen A Ercan, Carolyn M Hendriks, and Simon Niemeyer. 2017. Twelve key findings in deliberative democracy research. *Daedalus* 146, 3 (2017), 28–38.
- [8] Nicholas Diakopoulos. 2017. Enabling accountability of algorithmic media: transparency as a constructive and critical lens. In *Transparent Data Mining for Big and Small Data*, Tania Cerquitelli, Daniele Quercia, and Frank Pasquale (Eds.). Springer, 25–43.
- [9] Nicholas Diakopoulos and Michael Koliska. 2017. Algorithmic transparency in the news media. *Digital Journalism* 5, 7 (2017), 809–828.
- [10] Michael D Ekstrand, Daniel Kluver, F Maxwell Harper, and Joseph A Konstan. 2015. Letting users choose recommender algorithms: An experimental study. In *Proceedings of the 9th ACM Conference on Recommender Systems*. ACM, 11–18.
- [11] Michael D Ekstrand and Martijn C Willemsen. 2016. Behaviorism is not enough: better recommendations through listening to users. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 221–224.
- [12] Simone Fischer-Hübner, Julio Angulo, Farzaneh Karegar, and Tobias Pulls. 2016. Transparency, Privacy and Trust–Technology for Tracking and Controlling My Data Disclosures: Does This Work?. In *IFIP International Conference on Trust Management*. Springer, 3–14.
- [13] Gerhard Friedrich and Markus Zanker. 2011. A taxonomy for generating explanations in recommender systems. *AI Magazine* 32, 3 (2011), 90–98.
- [14] David Graus, Maya Sappelli, and Dung Manh Chu. 2018. Let me tell you who you are. In *Proceedings of the FATREC 2018 Workshop: Responsible Recommendation*.
- [15] Jaron Harambam, Natali Helberger, and Joris van Hoboken. 2018. Democratizing algorithmic news recommenders: how to materialize voice in a technologically saturated media ecosystem. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376, 2133 (2018), 20180088.
- [16] Jaron J Harambam. 2017. "The Truth Is Out There": Conspiracy culture in an age of epistemic instability. (2017).
- [17] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications* 56 (2016), 9–27.
- [18] Natali Helberger, Kari Karppinen, and Lucia D'Acunto. 2018. Exposure diversity as a design principle for recommender systems. *Information, Communication & Society* 21, 2 (2018), 191–207.
- [19] Dietmar Jannach, Sidra Naveed, and Michael Jugovac. 2016. User control in recommender systems: Overview and interaction challenges. In *International Conference on Electronic Commerce and Web Technologies*. Springer, 21–33.
- [20] Yucheng Jin, Nava Tintarev, and Katrien Verbert. 2018. Effects of personal characteristics on music recommender systems with different levels of controllability. In *Proceedings of the 12th ACM Conference on Recommender Systems*. ACM, 13–21.
- [21] René F Kizilcec. 2016. How much information?: Effects of transparency on trust in an algorithmic interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 2390–2395.
- [22] Bart P Knijnenburg, Saadhika Sivakumar, and Daricia Wilkinson. 2016. Recommender systems for self-actualization. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 11–14.
- [23] Jaya Kumar and Nava Tintarev. 2018. Using Visualizations to Encourage Blind-Spot Exploration. In *Recsys workshop on Interfaces and Decision Making in Recommender Systems*.
- [24] Stuart E Middleton, Nigel R Shadbolt, and David C De Roure. 2004. Ontological user profiling in recommender systems. *ACM Transactions on Information Systems (TOIS)* 22, 1 (2004), 54–88.
- [25] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 3111–3119.
- [26] Eli Pariser. 2011. *The filter bubble: What the Internet is hiding from you*. Penguin UK.
- [27] Denis Parra, Peter Brusilovsky, and Christoph Trattner. 2014. See what you want to see: visual user-driven approach for hybrid recommendation. In *Proceedings of the 19th international conference on Intelligent User Interfaces*. ACM, 235–240.
- [28] Frank Pasquale. 2015. *The black box society*. Harvard University Press.
- [29] Urbano Reviglio. 2019. Serendipity as an emerging design principle of the infosphere: challenges and opportunities. *Ethics and Information Technology* (2019), 1–16.
- [30] Yuri Saito and Takayuki Itoh. 2011. MusiCube: a visual music recommendation system featuring interactive evolutionary computing. In *Proceedings of the 2011 Visual Information Communication-International Symposium*. ACM, 5.
- [31] Rashmi Sinha and Kirsten Swearingen. 2002. The role of transparency in recommender systems. In *CHI'02 extended abstracts on Human factors in computing systems*. ACM, 830–831.
- [32] Maartje ter Hoeve, Mathieu Heruer, Daan Odijk, Anne Schuth, and Maarten de Rijke. 2017. Do News Consumers Want Explanations for Personalized News Rankings. In *FATREC Workshop on Responsible Recommendation Proceedings*.
- [33] Nava Tintarev. 2017. Presenting Challenging Recommendations: Making Diverse News Acceptable. In *FATREC Workshop on Responsible Recommendation at Recsys'17*.
- [34] Nava Tintarev and Judith Masthoff. 2006. Similarity for news recommender systems. In *Proceedings of the AH'06 Workshop on Recommender Systems and Intelligent User Interfaces*. Citeseer.
- [35] Nava Tintarev and Judith Masthoff. 2007. A survey of explanations in recommender systems. In *2007 IEEE 23rd international conference on data engineering workshop*. IEEE, 801–810.
- [36] Nava Tintarev and Judith Masthoff. 2011. Designing and evaluating explanations for recommender systems. In *Recommender systems handbook*. Springer, 479–510.
- [37] Nava Tintarev, Shahin Rostami, and Barry Smyth. 2018. Knowing the Unknown: Visualising Consumption Blind-Spots in Recommender System. In *ACM Symposium On Applied Computing (SAC)*.
- [38] Nava Tintarev, Shahin Rostami, and Barry Smyth. 2018. Knowing the Unknown: Visualising Consumption Blind-Spots in Recommender System. In *ACM Symposium On Applied Computing (SAC)*.
- [39] Nava Tintarev, Emily Sullivan, Dror Guldin, Sihang Qiu, and Daan Odijk. 2018. Same, Same, but Different: Algorithmic Diversification of Viewpoints in News. In *UMAP workshop on Fairness in User Modeling, Adaptation and Personalization, in association with UMAP'18*.
- [40] Edwin B Wilson. 1927. Probable inference, the law of succession, and statistical inference. *J. Amer. Statist. Assoc.* 22, 158 (1927), 209–212.
- [41] Malte Ziewitz. 2016. Governing algorithms: Myth, mess, and methods. *Science, Technology, & Human Values* 41, 1 (2016), 3–16.