FD Mediagroep: Opening the black box of user profiles in content-based recommender systems

Participants: David Graus (company representative), Feng Lu (company representative), Dung Manh Chu (company representative), Anca Dumitrache (company representative), Nava Tintarev (academic leader), Dimitrios Bountouridis, Jaron Harambam, Domokos Kelen, Felicia Löcherbach, Mykola Makhortykh, Shabnam Najafian, Emily Sullivan, Marieke Vinkenoog, Daricia Wilkinson.

Context and problem

Personalized experiences powered by recommender systems have, after years of being a mostly academic endeavor, finally permeated our daily lives. Whether it is through personalized recommendations in web shops (e.g., Amazon), personalized media consumption (e.g., Spotify, Netflix), search engines, or virtual assistants (e.g., Apple Siri, Amazon Echo). However, driven by data breach incidents and ad-driven business models, we've recently seen a rise in distrust and skepticism around the collection of personal data—a requirement for recommender systems. In addition, the GDPR has generated an increased interest in aspects of explainability and transparency of black-box machine learning algorithms and models.

In this project, we aimed to explore methods of explaining one aspect of how our content-based recommender system works: *the user profile*. More specifically, we aimed to automatically summarize and visualize the recommender system's high dimensional internal representations of users. These profiles are automatically constructed from their reading behavior, by leveraging attributes of items, e.g., topics, entities, and tags.

A multi-disciplinary team of 9 participants tackled this problem of generating explanations for news recommendations. The team consisted of PhD students and postdocs, and even attracted participants who travelled from Hungary and the USA for the event.

Research approach

We first developed a conceptual framework where we identified different levels of explanations, and different values that users and FD media may share.

Following exploratory data analysis, we developed mockups for the interface for different explanatory values. This lead to the development of an interactive explanation interface, as well as a static version for which we planned a user-study.

Results

The team first identified three levels for explanations:

1. **Transparency.** "What is my interaction with the system so I can understand the past?" Explanations on the individual level. E.g., histograms of topics previously read.
2. **Contextualization. "**How does my interaction compare with my community, so I can understand myself within this context?" Explanations relating individual consumption to

aggregated consumption. This comparison could be with the general population, public history, or driven by company interests.

3. **Self-actualization.** "How can I use the system so I can reach my epistemic goals?" This relates to explanations relating to the values shared both by users and news providers. We identified four values: *broaden horizons, discover the unexplored, expertise, and stay informed.*

To focus the work, we centered our work on two of these values:

Broaden Horizons

· **How it fits with FD**: Presenting a broad and balanced perspective on what matters

· **Explanation for the user:** ...because you want to see many angles, and are interested in inclusivity.

Discover the Unexplored

· **How it fits with FD**: broadness in terms of promoting marginalized voices to remain unbiased

· **Explanation for the user**: ...because getting a new perspective promotes objectivity

Broaden horizons focuses on gradual expansion of tastes: topic diversity, filling nearby gaps, and incremental diversity. In contrast, *Discovering the unexplored* is focused on getting users to discover (completely different) topics they haven't yet read about, long-tail diversity, and foster exploration. Broaden horizons is in other words more gradual. For any solution, we also identified that it is equally important to support informed abandonment of topics (choosing what not to read) as well as supporting acceptance of topics.

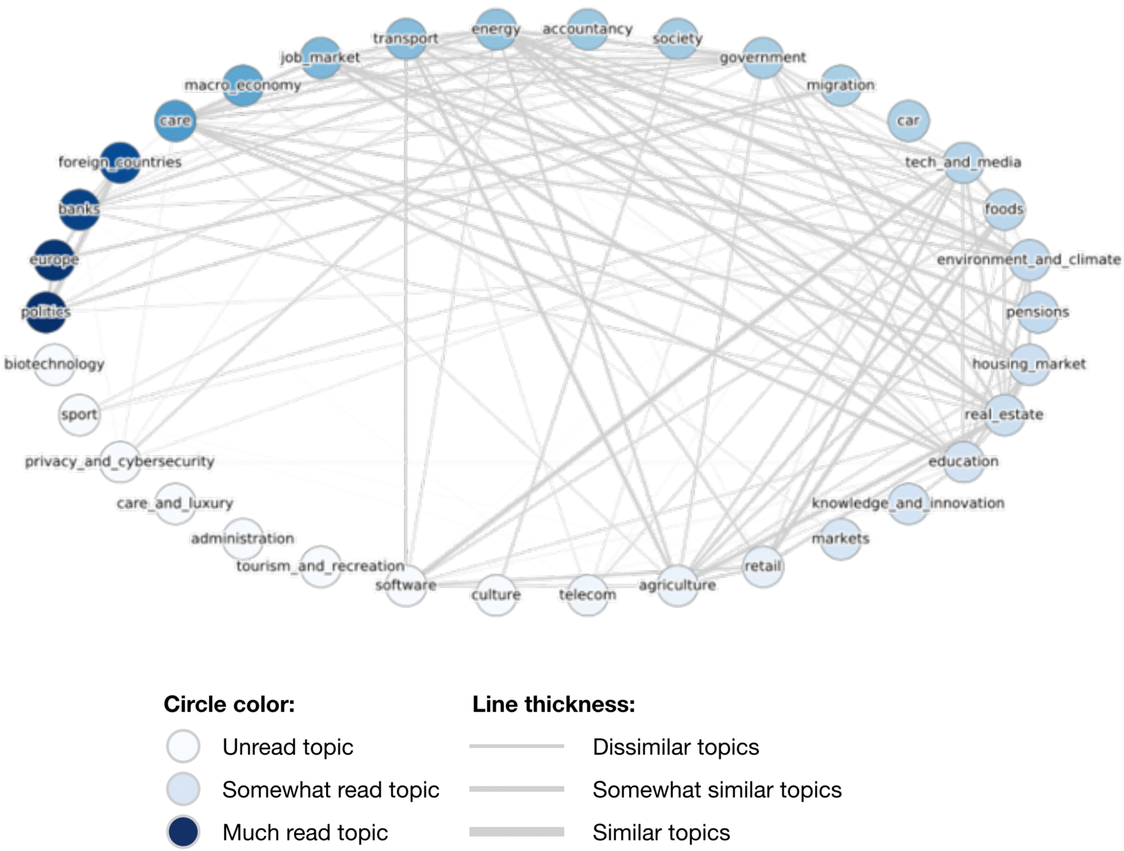

We developed a visualization (see Figure above) which compares a user profile with a collection of articles (using approximately 1600 articles released by FD, which represents the articles read by a sample of users during November 2018). This profile allows a user to view their reading habits with the content released by the newspaper.

A darker color (e.g., for *Politics and Europe*) reflects topics the user has read more often, and a lighter color (e.g., *Biotechnology and Sport*) reflects a topic the user has consumed less often or not at all. The topics are shown decreasing order (most read first). This aspect of the visualization allows users to identify topics that they read often, and where they may have blindspots (topics they have not read anything about).

Furthermore, connections are made between topics using chords. A thicker chord reflects a higher similarity between the topics based on their co-occurrence across all articles (using word embedding). For example, this model finds that the topic *Politics* is very similar to *Foreign Affairs* and *Europe.*

We planned and ran a crowd-sourced user study for a static version of this visualization. This study measured both the effectiveness of the visualization and its understandability. In a within-subjects design, we compared the behavior of users for the two values (Discover the unexplored, and Broaden horizons). This study used 4 anonymized user profiles from existing users of FD. To improve the realism of the experimental task, users were asked to identify which they felt best described their interests.

Our first results suggest that users explore more topics they have not read before in the Discover the unexplored condition, compared to the Broaden horizons condition. However, we require more time to collect sufficient data and complete the analysis to check whether this result is statistically significant.

An interactive version of the visualization can also be found online here: http://info.ilab.sztaki.hu/~kdomokos/ict/

Future work

In our further work we plan to apply the framework to develop user studies for the other value-driven goals (i.e.., Expertise and Staying informed).

Another idea we explored at the workshop was hierarchical visual representation of topics. While this hierarchical information was not well represented in the original data, we identified a visualization that would be suitable for this, and requires further evaluation.

We also identified extensions to the visualization that were omitted to focus the evaluation on a limited number of parameters. One of these consisted of a further development of level two transparency: comparing a user to their community e.g., by comparing a user profile to other users or the catalog. The developed visualizations are suitable for this kind of expansion, e.g., by graying out parts of the interface explored by other users but that are underrepresented by the current user. These were not included in the user-study however to ensure that we were evaluating the effect of the value (broaden horizon, or discover the unexplored) rather than a large number of interface elements.

We are excited by all the ideas and designs that developed during the course of the workshop. We plan a joint publication, and further exploring the topic in a continued collaboration!