Appendix A: Additional Usage Scenario

Escaping a Potential Filter Bubble

In this scenario, the user sees the difference in recommendations between two profiles and by tweaking the configuration, they are able to make the recommendations more similar, reducing the effect of a filter bubble. The user starts out by selecting the profile that matches their interests best, in the case of this scenario: the politician. In another instance of the workspace, the user selects another profile, the athlete. The scatter plots depicted in Figure 3 show that the recommended articles come from different clusters. Furthermore, the user observes in the graph (Figure 4) that they are only receiving articles recommended from three categories, sports, entertainment and TV, hinting at a possible filter bubble. The user utilizes the graphs in Figure 3 and Figure 4 to assess the cause of this filter bubble and changes the model configuration in Figure 2. The user increases the contribution of popularity in Figure 2a, reduces the number of articles considered in Figure 2b, and increases exploration in Figure 2c. Following this, the user identifies a diminishing impact of the filter bubble. However, the user is not as satisfied with these recommendations and reverts the changes back as he preferred those recommendations. The user experiences empowerment, as they are in control of their filter bubble and always have the opportunity to escape if preferred. This empowerment enables the user to determine or reduce their personal filter bubble. The user starts by selecting two profiles and comparing the initial recommendations. Then the user tweaks the parameters using the tools demonstrated in Figure 2 and compares the outputs again. This way, the user profile's and algorithm's impacts become visible over multiple iterations. This process can generate trust, transparency, and understanding, which will also translate to other news RSs outside this workspace. For example, the user could be more aware of their online click behavior because they learned how much impact this has via the user profile on their recommended news.

Appendix B: MiND statistics

Figure 5 shows the distribution of articles in the MiND dataset (see Subsection 3.2). The most common categories are news and sports, containing more than 10,000 articles. Other categories have fewer than 5,000 articles.

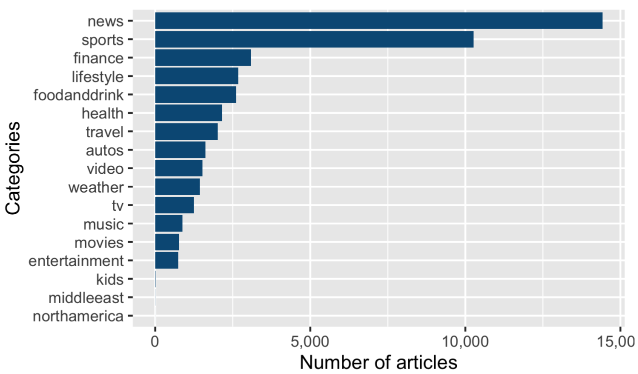

Figure 5: Distribution of articles between categories