The Recommendation Dashboard:
A System to Visualise and Organise Recommendations

Gerwald Tschinkel, Cecilia di Sciascio, Belgin Mutlu and Vedran Sabol

Know-Center, Graz University of Technology
gtschinkel, cdisciascio, bmutlu, vsabol@know-center.at

Abstract

Recommender systems are becoming common tools supporting automatic, context-based retrieval of resources. When the number of retrieved resources grows large visual tools are required that leverage the capacity of human vision to analyse large amounts of information. We introduce a Web-based visual tool for exploring and organising recommendations retrieved from multiple sources along dimensions relevant to cultural heritage and educational context. Our tool provides several views supporting filtering in the result set and integrates a bookmarking system for organising relevant resources into topic collections. Building upon these features we envision a system which derives user's interests from performed actions and uses this information to support the recommendation process. We also report on results of the performed usability evaluation and derive directions for further development.

1 Introduction

Within the EEXCESS research project[1] work is performed to increase the visibility of cultural, scientific, and educational content by injecting automatically generated recommendations from various information sources into the Web pages visited by the user. In this context, recommender systems are used for pro-active, context-based retrieval of potentially relevant resources. However, recommender systems leave little or no control to the user over the actually performed search, potentially resulting in a set of recommendations covering a broader set of topics with the number of retrieved resources being too large for browsing and individual inspection. This necessitates development of tools, which support exploration and filtering in the retrieved recommendation set, and provide means for collection of resources satisfying particular conditions.

In this work, we describe our work in progress on a visual system for exploring and filtering in recommender

results along dimensions relevant to cultural and educational context, primarily temporal and geographical information, but also along metadata such as language, source or mime type. Additionally, our system provides an integrated bookmarking subsystem allowing users to organise recommendations of interest into collections, whereby the collections themselves can be visualised and analysed further using our tool. With the goal of improving the recommendation quality, we perform in-browser extraction of keywords and dominant topics from the user created collections. We are working on automatic integration of the extracted information into the "Topics of Interest" section of the user-profile (which is currently edited manually) to increase the relevance of the retrieved recommendations.

We also introduce the results of a preliminary usability evaluation performed using the thinking aloud method. We report on the collected user feedback, describe new features currently under development and provide plans for future improvements of the design.

2 Related Work

Richard et al.[10] underline how important it is to decrease the barriers for accessing cultural information by automating the creation of dashboards, a metaphor having its origin in the industry and economics [2], where it is used to give an overview of different data aspects. Recommender systems provide a variety of metadata facets, which can be presented in the visualisations of a dashboard. The presentation is known to significantly influence user experience and effectiveness of the recommendations [5].

For presentation of search results [6] discuss visual alternatives to the ranked list advocating the use of visualisation for aiding discovery. Text-based faceted search interfaces, introduced by [3], provide an effective way for filtering and drilling down in result sets along metadata dimensions. More recent work includes FacetMap [9] which

---

[1] http://eexcess.eu

visualises different facets as oval bubbles grouped and coloured depending on the facet type, FacetAtlas [1] that shows facets as a graph in a multi-facet relational diagram, and FacetScape [8] that displays facets as tag clouds within Voronoi cells representing faceted categories. These tools are visualising facets in a generic way, without considering specific data characteristics distinguishing each facet category. APA Labs [4], in contrast, employs specific visualisations for each facet (e.g. timeline or geo-visualisation).

In contrast to the above listed systems our Recommendation Dashboard (RD) provides both the facet-specific visualisations supporting exploration and filtering of retrieved results, and the possibility to organise interesting result into collections which can be explored further using the provided visual means. In addition, RD is designed to support the recommendation process by inferring user's interests from the collected recommendations and from the interactions performed while exploring the data.

3 Visualisation of Recommender Results

Our usage scenario focuses on exploration, filtering and organisation of recommender results. Figure 1 shows a high level view of the employed system architecture. The scenario starts with the user viewing or composing content which is captured by a Browser Extension component (currently available only for the Chrome Browser). It injects background scripts into every visited Web page to access the DOM-tree, retrieve the content and compute a summarised representation of the page (keywords) which is sent to the federated recommender.
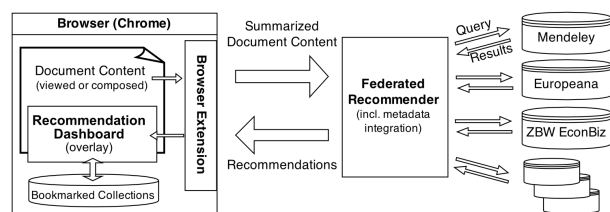

Figure 1: System architecture comprises a Browser Extension, a recommender, several data sources and the Recommendation Dashboard tool.

The recommender retrieves results from various data sources, e.g. Europeana for cultural heritage resources or Mendeley for scientific papers. It integrates the metadata formats and returns the recommendations back to the Browser Extension, which displays the results in the form of a list in a sidebar integrated in the right side of the currently viewed Web page. The Browser Extension integrates visual tools which are activated with a click of a button and displayed as an overlay to the viewed page. In the following we describe the RD, a Web-based visual tool for exploring and organising the retrieved recommendations.

3.1 Recommendation Dashboard

As the size of the recommendation space grows to tens or hundreds of items, it becomes increasingly difficult for the user to explore and analyse them and to pick those which are actually interesting. The RD leverages the capacity of the human visual system for dealing with large amounts of information at once by providing a multi-visualisation framework that allow users to explore, filter and organise recommendations in the area of cultural heritage and education according to features that typically characterise them. RD primarily targets pupils, researchers and the general public, and should be easy enough to use for people without experience with visual data analysis tools. For this reason we provide methods for automatic binding of visualisations to the data and refrain from using complex techniques such as multiple coordinated visualisations.

The main view of the user interface is subdivided in five areas according to their roles (see Figure 2). At the top left, the Query Panel provides the possibility to enter and execute a search query. On the top right in the Dataset Panel the user can choose whether the currently retrieved set of recommender results or one of the bookmarked collections are shown in the RD. The List on the right shows the actual resources where each entry includes the title, source, language, a thumbnail (if available) and a possibility to bookmark the resource. The visualisation in the middle can be chosen in the Control Panel, located on the left. There, the visualisation can also be configured by manipulating the mapping of different metadata facets onto the visualisation's visual properties, such as axes or colours. Only meaningful mappings from the data onto the visual properties are permitted, e.g. categorical values can be mapped onto colours or onto an axis, while numeric values can only be mapped onto an axis. We compute the mappings by matching the data characteristics with semantic descriptions of visualisation's visual properties, as described in [7].

The metadata describing recommendations in the cultural heritage domain are primarily time stamps, geo-references, language, data provider and mime-type information. Therefore, the RD currently offers the possibility to choose from one of three different visualisations, namely a Timeline, a GeoView, and a BarChart, each one of them supporting faceted exploration from a different perspective (see Figure 2). All charts incorporate features aiming at facilitating faceted exploration. In the **Timeline** the x-axis is always bound to the date while the y-axis and the colour can be mapped to categorical faceted attributes. The chart incorporates a temporal brush. **The GeoView** shows the
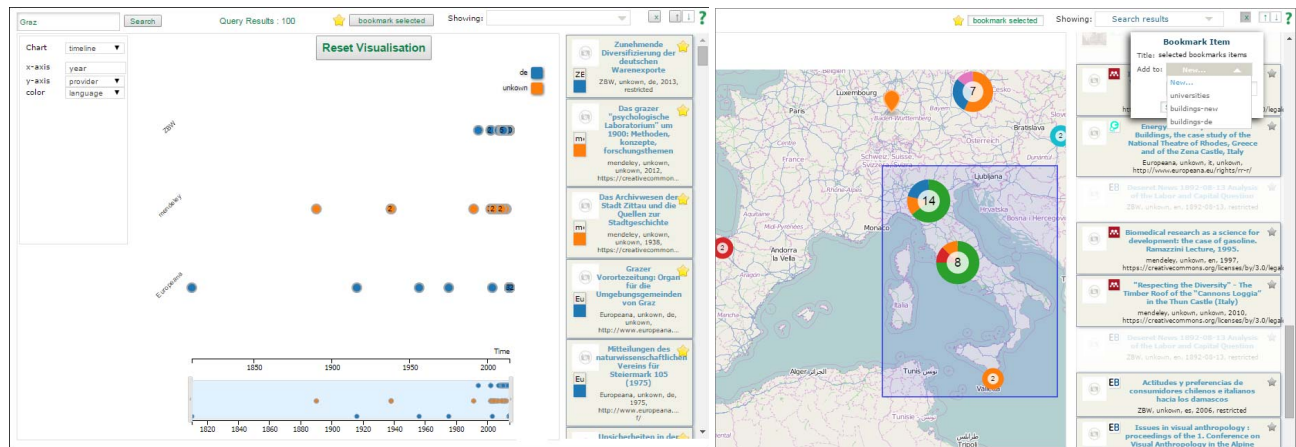
Figure 2: Visualisation Dashboard showing recommendations. The Timeline is selected for visualisation (left). Bookmarking of multiple selected items on the geo-visualisation or of single items in the list (on right).

distribution of geo-references included in the recommendations. When the density of single data items on screen becomes too high for displaying them separately they will be aggregated into a single symbol – a "donut chart" – representing the distribution of the categorical metadata and displaying the number of the underlying items in the centre. The chart includes a rectangular brush tool. **The Bar-Chart** is used to represent categorical attributes (e.g. language or data provider) on the x-axis and a numerical attribute on the y-axis. Filtering is performed by clicking on the bars. A screencast showing the basic usage of the system, is available at [2].

3.2 Features Under Development

A major feature we are working on, addresses improving the recommendation quality by deriving user interests automatically and incorporating these into the user profile. For this purpose we use an in-browser text analysis engine (based on natural.js) to compute keywords from the recommendations belonging to the user's collections. The highest frequency keywords of a selected collection will be automatically added to topics of interests within the users profile to improve the relevance of the retrieved results.

Also in development is an information landscape visualisation that focuses on topical exploration. The landscape shows major topical clusters described by keywords which can be selected by the user to filter the recommendations. In doing so the user expresses interests which, again, shall be used for steering the recommendation process.

4 Evaluation

The primary question to be answered by the usability evaluation was whether and how users would be able to

utilise the built-in functionality and where the major issues were located. Additionally, we wanted to find out which additional functionality the users would consider useful. The tasks where only about the user experience of the dashboard including its functionalities but not about the recommender results. The evaluation was performed on the stable feature set as described in Section 3.1.

4.1 Methodology

The study was conducted using the thinking aloud method, where users are asked to comment freely and express their thoughts while performing tasks. Everything participants said and did was noted. We began with a free trial of the tool to let users familiarise with the system. Then the participants executed a series of tasks (time limit 5 minutes per task):

Task 1: Selecting recommendations geographically
Task 2: Temporal selection of recommendations
Task 3: Filtering by categorical metadata
Task 4: Data cleansing depending on the covered topics

After performing the tasks users were asked to provide comments and to suggest additional functionality.

Data: As the returned recommendations may vary over time, we ensured that all users test with the same data by preparing two test data sets. The smaller one on "University Campus" (36 recommendations) was used for the trial of the tool. The larger one on "Historic Buildings" (68 recommendations) was used for the test tasks.

Participants: 12 users participated in the study, aged between 18 and 38, without colour blindness or other relevant impairments. 10 people had a background in IT.

---

[2] https://www.youtube.com/watch?v=IKwEexY_Ay4

4.2 Results

We have combined the feedback collected from thinking aloud with the written comments collected in the questionnaires. We grouped the 210 collected comments (17.5 per participant in average) into 78 issues. Next, we assigned each issue to one or more of the 10 high-level categories as shown in Figure 3. In the following we briefly outline the major issues found.
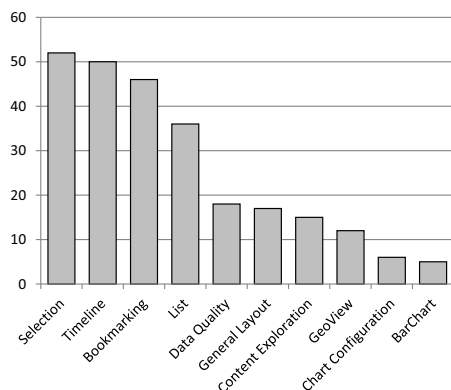

Figure 3: Number of issue mentions per category.

Selection mechanisms provided by the employed visualisation libraries (D3.js, Leaflet) were not sufficient. Mechanisms enabling very precise, adjustable selection (supported by zooming) were requested, e.g. to pick locations close to a country border or time-stamps with high accuracy over large time intervals. Several components suffered from usability issues which are comparatively easy to address, in particular the layout problems of the **Timeline** (overlapping items although enough free area available), the **Bookmarking** in having difficulties locating the right functionality in the UI and requiring simplified dialogues, and the **List** with the inability to control the sorting and the confusing auto-scroll on selection behaviour. The **Data Quality** turned out to be a major issue with participants primarily expressing concerns over the trustworthiness of the resources, but also criticising unclear metadata qualifiers, and questioning whether the right recommendations have been retrieved. Clearly, the users wanted regain at least some control over the recommendation process. Finally, for **Content Exploration** better tool support was requested (e.g. in page search or content summaries).

5 Conclusion and Future Work

We have introduced the Recommendation Dashboard, a tool for exploration and organisation of recommender results retrieved from multiple sources of cultural, educational and scientific content. Evaluation results indicate that facet-specific visualisations were very well-received and that the design of the current version of the tool is adequate for its intended tasks. The results also point to directions for further improvements and new functionality. We are currently working on integrating tools for content-based exploration and on improvement of recommender performance through automatically extending the user profile with user's interests derived from the bookmarked collections. We also started work on integrating multimedia content (i.e. images) into the existing visualisations, implemented as cover-flow as a remedy to overlap issues.

In the future we will add more powerful selection mechanisms, and provide support for multiple active filters represented by simple micro-visualisations, which should keep the visual complexity of the UI low while providing functionality similar to Multiple Coordinated Views. We also plan to explore visual means for conveying trust-enabling facets. By working on the identified issues and by extending the Recommendation Dashboard we are looking forward to deploying the software to a wider audience by the end of 2015 where we will conduct additional users feedback via a/b testing and a larger usability evaluation.

References

[1] Nan Cao, Jimeng Sun, Yu-Ru Lin, David Gotz, Shixia Liu, and Huamin Qu. Facetatlas: Multifaceted visualization for rich text corpora. *IEEE Trans. Vis. Comput. Graph.*, 2010.

[2] Stephen Few. *Information Dashboard Design: The Effective Visual Communication of Data*. O'Reilly Media, Inc., 2006.

[3] M. A. Hearst. Design recommendations for hierarchical faceted search interfaces. In *Proc. SIGIR 2006, Workshop on Faceted Search*, 2006.

[4] Wolfgang Kienreich, Elisabeth Lex, and Christin Seifert. APA labs: An experimental web-based platform for the retrieval and analysis of news articles. IEEE Computer Society, 2008.

[5] Joseph A. Konstan and John Riedl. Recommender systems: From algorithms to user experience. *User Modeling and User-Adapted Interaction*, 2012.

[6] William Kules, Max L. Wilson, m.c. schraefel, and Ben Shneiderman. From keyword search to exploration: How result visualization aids discovery on the web. Technical report, 2008.

[7] Belgin Mutlu, Patrick Hoefler, Gerwald Tschinkel, Eduardo Veas, Vedran Sabol, Florian Stegmaier, and Michael Granitzer. Suggesting Visualisations for Published Data. In *Proceedings of IVAPP 2014*, 2014.

[8] Christin Seifert, Johannes Jurgovsky, and Michael Granitzer. Facetscape: A visualization for exploring the search space. In *Proceedings of Information Visualisation 2014*, 2014.

[9] Greg Smith, Mary Czerwinski, Brian Meyers, Daniel C. Robbins, George G. Robertson, and Desney S. Tan. Facetmap: A scalable search and browse visualization. *IEEE Trans. Vis. Comput. Graph.*, 2006.

[10] Richard J. Urban, Piotr Adamczyk, and Michael B. Twidale. Cultural Heritage Information Dashboards. In *Proceedings of iConference 2010*, 2010.