Using Micro-Visualisations to Support Faceted Filtering of Recommender Results

Gerwald Tschinkel, Robert Hafner, Peter Hasitschka, and Vedran Sabol

Know-Center GmbH, Graz University of Technology
gtschinkel@know-center.at, robert.hafner@student.tugraz.at, hasitschka@student.tugraz.at, vsabol@know-center.at

Abstract

Faceted search is a well known and broadly implemented paradigm for filtering information with various types of structured information. In this paper we introduce a multiple-view faceted interface, consisting of one main visualisation for exploring the data and multiple miniaturised visualisations showing the filters. The Recommendation Dashboard tool provides several interactive visualisations for analysing recommender results along various faceted dimensions specific to cultural heritage and scientific content. As our aim is to reduce the user load and optimise the use of screen area, we permit only one main visualisation to be visible at a time, and introduce the concept of micro-visualisations – small, simplified views conveying only the necessary information – to provide natural, easy to understand representation of the the active filter set.

Keywords— **Recommendation Visualisation, Faceted Filtering, Multiple Views, Micro-Visualisations**

1 Introduction

The aim of our research was to support users in finding relevant information within the space of cultural heritage and science. The EEXCESS Project[1] focuses on recommending content to the user by discovering what the user is interested in while surfing the web, and injecting recommended documents into a popover control. As the number of recommended documents grows, it becomes more and more difficult to analyse and explore them, making it essential to provide a visual tool that leverages the users capabilities. The Recommendation Dashboard provides several interactive visualisations, which allow the user to analyse recommendations in different perspectives. Brushing is supported along all provided metadata dimensions, with multiple filter functionality supporting data drill-down and removing of non-relevant recommendations. As the metadata dimensions differ in type (spatial, temporal, categorical, textual etc.) and in meaning, different representations are required for different filters. The user needs a way to visualise all filters applied to the actual data set, however without being overloaded with the complexity of a multi-visualisation interface.

This paper describes the latest version of the *Recommendation Dashboard* (RD) interface which aims at ameliorating the situation: we introduce micro-visualisations which display just enough information to convey the filters and data ranges which are currently applied. Each active filter is represented by a simplified, data type-specific visual representation, which is designed to reduce the load on the user, and at the same time provide a more natural representation for the particular information type then a standard text-based faceted interface.

2 Related Work

In many cases, human perception is limited by the large number of results delivered by web search engines or recommender systems. It is a great inconvenience for the user to look through a list containing many items in order to find the ones which are really needed. The metadata of the search results can be used for filtering and reducing the size of the result set. This approach is known as faceted search and is a proven technique for exploration of complex search result sets [6] using multiple filters which are set in different meta data categories. However, a typical faceted search user interface is text based, which means that various meta data types (e.g. temporal, spatial or categorical) are not displayed in the most suitable visual form. A more advanced form is provided by the FaceScape visualisation [15] where metadata categories are visualised as areas of a Voronoi diagram, with facet filters shown as

[1] http://www.eexcess.eu

2375-0138/16 $31.00 © 2016 IEEE
DOI 10.1109/IV.2016.77

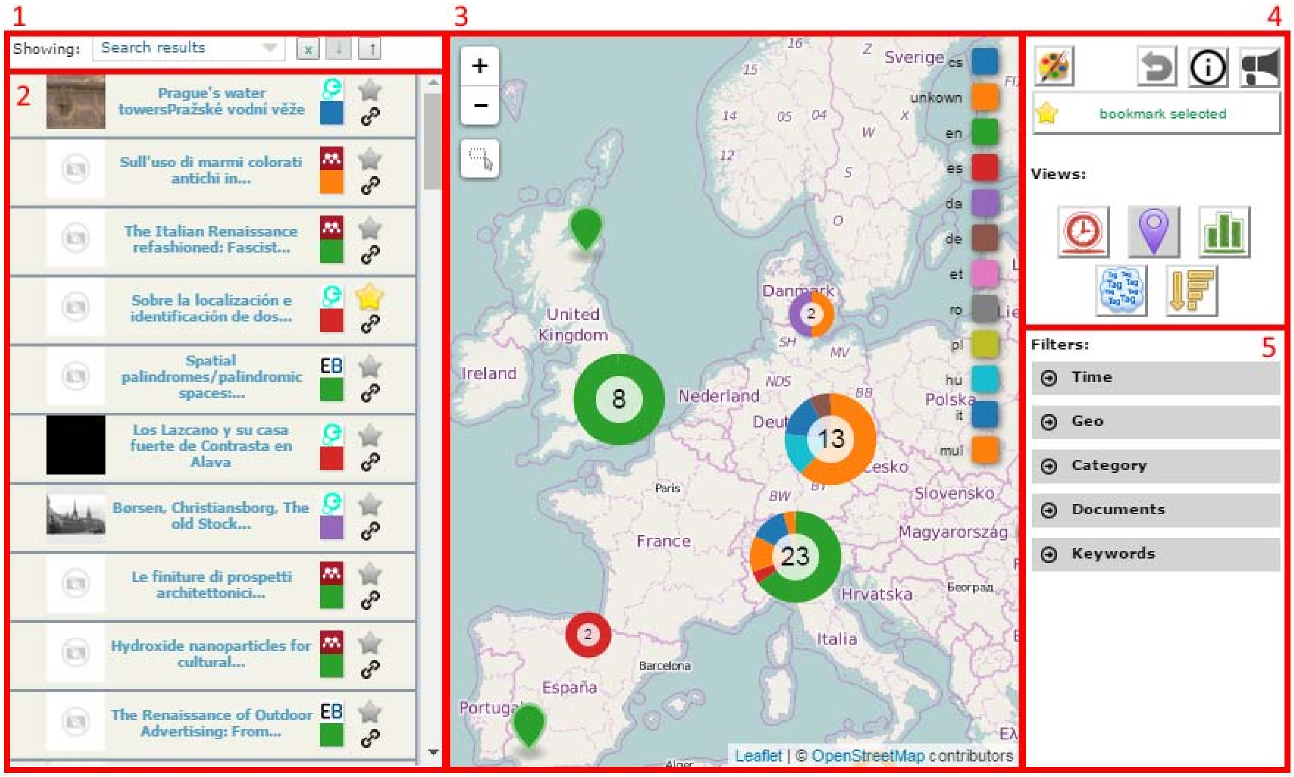

Figure 1: The Recommendation Dashboard user interface showing a set of recommendations related to "renaissance". A geo-visualisation is selected as a main view showing the distribution of geo-reference metadata.

tag clouds residing within the corresponding areas. While information such as category size or facet weight are represented by visual properties, textual representation remains the main carrier of information independently of the data type. Faceted browsing has found wide use in various areas, including information retrieval and commercial search applications [11], biology [4], medicine applications, learning [5], academic search [1] and eCommerce applications (e.g. Amazon). The design of such a system poses a major challenge. The user acceptance and satisfaction is known to depend on the quality of the retrieval and the usability of the application [2]. In [9] authors define important objectives for achieving a satisfying result by evaluating the Flamenco search interface project[2].

Multiview visualisation is a proven concept in information visualisation for dealing with heterogeneous data sets [12]. Various aspects of the data are visualised using multiple specialised views, typically between two and four.[10]. An example of a multiple view interface for Web-based search is VisGets [3], where temporal, geo-spatial, and topical filters can be set using the corresponding visualisations. However, the authors of such systems do not consider the increased load on the user which is introduced by multiple complex views which are visible at once. In the ApaLabs [8] only one of the several available visualisations is visible at any time. Each visualisation can be used for filtering, but multiple active filters are not supported. An early RD version [17] followed an approach similar to ApaLabs. An initial evaluation revealed that even a single interactive visualisation introduces a considerable load on the typical user, but nevertheless multiple filter support ranked high on the users' wish list. We propose a solution where only one out of several interactive visualisations is visible at each moment. To reduce complexity, the active filters are represented by a set of simplified, data type-specific views - the micro-visualisations.

3 Recommendation Dashboard

While the user browses the Web, a recommender engine can automatically suggest a lot of potentially useful resources. However, as user's interests and current needs are not easy to infer, a significant proportion of the retrieved resources may not be relevant. Interactive visual tools are applied to address this issue and provide the users

[2]http://flamenco.berkeley.edu

with a way to gain an overview of the results, refine their interests and communicate these to the system.

Our usage scenario focuses on the exploration, filtering and organisation of recommender results. The scenario starts with the user viewing or composing content, which is captured by a Browser Extension component [13] (currently available only for the Chrome Browser). The Browser Extension injects background scripts into every Web page visited to access the DOM-tree, retrieves the content and computes a summarized representation of the page (keywords and named entities), which is sent to the recommender. The federated recomender [7] retrieves and ranks resources from sources such as Europeana[3], ZBW EconBiz[4] or Mendeley[5]. The retrieved results are broadcast to various client components, including the Recommendation Dashboard (RD) interface [17].

The Recommendation Dashboard (Figure 1) supports users in exploring a set of recommendations retrieved by the recommender. It provides several interactive visualisations supporting analysis and filtering of recommendations along different perspectives specific to cultural heritage and scientific content. The features describing recommendations include both the faceted metadata (e.g. language, year, data source or geo-references) as well as content-based features (keywords and named entities). Single recommendations or filtered subsets of the result set can be bookmarked by assigning them to one of the user-defined collections. As the user's collections are saved in the browser, collaborative bookmarking is currently supported only by exporting and exchanging collections. A server-side collection storage is planned for a future RD version.

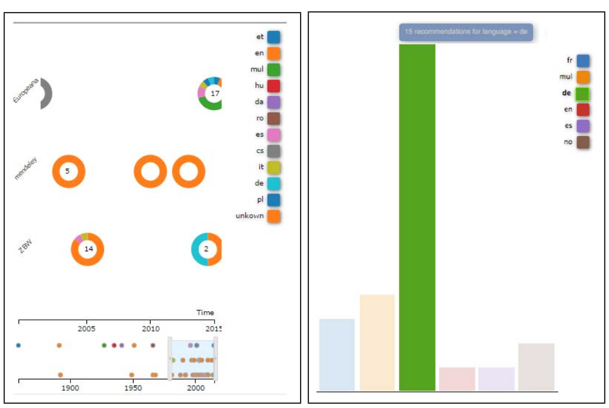

Figure 2: Timeline (left) with the temporal brush set, and a bar chart (right) showing languages with German (in green) selected as brush.

[3]http://www.europeana.eu/
[4]http://www.econbiz.de/
[5]https://www.mendeley.com/

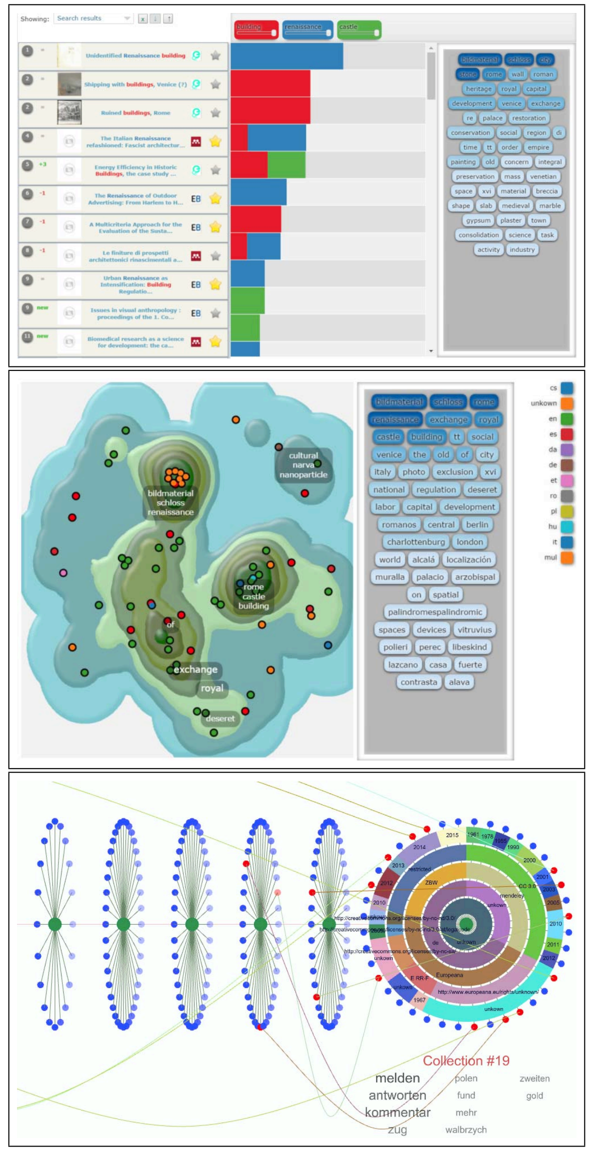

Figure 3: New RD main visualisations: uRank (top) and topical landscape (center) for topical exploration, and the history graph (bottom) for query history navigation.

In the latest version, the RD user interface (Figure 1) is divided into the following parts:

1. Data set area: lets the user choose the currently shown data set - either the recommender results (default) or one of the bookmarked collections.
2. Result List: shows the list of recommendations in the order returned by the recommeder. A recommendation is represented by the title, a thumbnail, an icon

representing the source, a color representing the language, a link button to insert the resource URL into an edited document (editing scenario only), and a bookmarking button (star) to add a resource to one of the user-defined collections.

3. Main Visualisation: shows one of the six interactive visualisations for exploring the recommendation set. Each visualisation supports brushing and navigation.
4. Control Elements: provide the possibility to change the main visualisation, to assign the visual property of color to a metadata category, and to bookmark the currently filtered subset of the data.
5. Filters Area: contains the micro-visualisations representing the active set of filters (see Section 4).

The architecture of the RD is extensible allowing for easy integration of additional main visualisations.Besides the timeline (Figure 2, left), the bar chart (Figure 2, right), and the geo-visualisation (Figure 1, center), which were initially available in the RD, the latest RD version now includes functionality requested during user testing [17]: i) provide transparency on how specifying topical interests affects the ranking of recommendations, ii) provide an overview of dominant topics covered by the result set, and iii) provide support for query history navigation.

The new visualisations supporting the requested functionality are shown in Figure 3. uRank (top) [14] gives an overview of the topics covered by the recommendations in the form of a tag cloud. The user selects interesting concepts and specifies their importance by adjusting weights with the corresponding sliders. The recommendations are re-ranked depending on user's choices and visualised using stacked bars encoding topical composition of each entry. Topical landscape (center) [16] places similar items close to each other resulting in dominant topical clusters being shown as groups of items. Clusters are visualised as islands labelled by important terms from the underlying documents. In both uRank and topical landscape selected keywords are used as topical brushes. The history graph (bottom) shows the query history, where each query (green oval) is placed on a timeline, while the corresponding recommendations are shown as a surrounding ring of items. Recommendations occurring in multiple result sets are connected by links. By clicking on a query the recommendation ring is rotated towards the user and the space within it is filled with the sunburst visualisation showing metadata distribution.

4 Micro-Visualisations

The brushing interaction takes place in the main visualisation to select a particular subset of the data. After a brush is set, a micro-visualisation shows up in the filter area to represent the brush. Recommendations which are not selected by the brush are shown semi-transparent in the list to convey the effects of the brushing. On top of each micro-visualisation two buttons are shown: a trashcan for removing the brush (and the corresponding micro-visualisation) and a lock for adding the current brush to the permanent filter set. If the user changes the main visualisation at this point, the brush and the corresponding micro-visualisation will be discarded. When a filter is added using the lock-button, all recommendations which are not within the filter range are removed from the list and from all visualisations. By combining different filters the user can quickly narrow down the data set to find the most useful resources. Should the need arise, each permanent filter can be removed individually with the trashcan button.

Figure 4: Filters shown using micro-visualisations (left). Equivalent text-based faceted descriptors are shown for comparison (right).

Currently five different micro-visualisations are provided (Figure 4, left): one for each type of facet currently returned by the content providers – temporal, geo-spatial, categorical, and topical – and an additional one showing selections of single recommendations. The visual design is data type-specific to convey the filter information in an adequate, natural manner. As we focus on reducing the user load and minimising the use of screen area, the design includes just enough information to represent the filters. To support pre-attentive processing, we always favour a visual form over text, which is used only where necessary. To compare the visual filter representation with the classical text-based faceted interface we also provide a corresponding textual representation (Figure 4, right). The visual representation conveys richer information while at the same time reducing the amount of information that needs to be read and interpreted by the user. Note that exploring certain facets in textual form may also require scrolling.

Time Stamps

The visualisation of time-oriented data uses a hexagon-shaped stacked timeline (Figure 5), which addresses the requirement of visualising all time references in a single compressed representation. The visualisation transforms the recommendations into ideograms, where a single object is represented by a square rotated by 90 degrees. The cumulative visual items are hexagonal shapes, which are horizontally divided depending on the relative share of the articles (see the right side of the Figure), enabling the the user to see the distribution of categorical metadata. The orientation of the axes with their accompanying information remain as in the main visualisation, whereby the data source descriptions are replaced by more compact icons.

We also experimented with a dynamic, non-linear time representation, where large gaps between pairs of time references are compressed leaving more space for shapes representing the recommendations. The weakness of this approach is the likely confusion caused by the non-linear width of the gaps, which we tried to solve by marking the gaps on the time line and encoding their width by shades of gray. However, initial user feedback led us to revert to a linear time scale, until a more intuitive design is found.

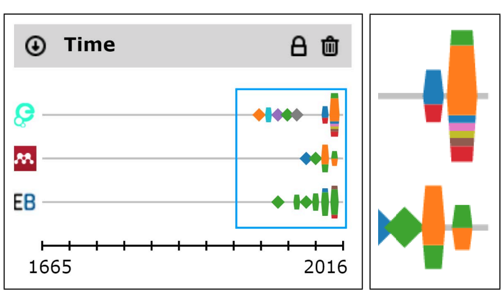

Figure 5: Time filter (left) with the detailed structure shown enlarged (right).

[6]http://www.bitmedia.at/
[7]http://www.zbw.eu/
[8]http://purl.org/EEXCESS/clients/chrome-extension
[9]http://eexcess.github.io/visualization-widgets-v3/examples/index-dashboard.html

Geo-spatial References

The main geo-visualisation supports a rectangular brush (represented by WGS 84 coordinates). The micro geo-visualisation displays the filter as a rectangle over the continents which are visually separated by using different shades of gray. A zoom-in to country level is supported to provide more detail when the filter covers a smaller area.

Categorical Metadata

Distribution of recommendations over categorical descriptors is visualised by a bar chart as the main visualisation. In contrast to that, the space-optimised micro-visualisation consists of filled hexagons placed next and above each other to minimise the waste of space. Hexagons are used instead of rectangles for two reasons: i) stacked rectangles could easily be mistaken for stacked bar charts, and ii) hexagons leave slightly more space in the middle for the label describing the category.

Topical Descriptors

The RD offers two main visualisations that provide topical filtering using keywords: uRank and topical landscape. In contrast to the discussed graphical micro-visualisations keywords can only be represented by text. In the filter area, the keywords are shown as a micro tag-cloud where shades of grey represent their weights (i.e. the number of documents containing the corresponding keyword).

Single Recommendation Selection

In the RD, the user can select one or more individual recommendations by clicking on them with the commonly used shortcut *Ctrl* + *mouseclick*. Each recommendation is represented by its thumbnail, with the full title being shown only on mouse-over inspection.

Implementation and Demos

The complete RD is implemented using Web technologies, namely HTML5, JavaScript and SVG, running in an enclosed IFRAME environment. The recommendations are injected via window messaging from an outside application, such as the Browser Extension. In addition to content consumption applications, the RD was integrated into two server-side content creation applications (with Bitmedia[6] and ZBW[7]) and is currently undergoing tests. The Browser Extension[8] containing the RD can be installed from the Chrome web store. A stand alone RD version[9] is also available and can be tried out with built-in test data.

5 Conclusions and Future Work

In this work we have introduced a new version of the RD tool for exploring and analysing recommender results. The tool was extended with new visualisations supporting topical exploration and history browsing. Most importantly, it now provides a novel visual representation for faceted filters using simplified, data type-specific representations - the micro-visualisations. They are designed to provide an overview of the current filter set in a compact, natural form which aims to minimise the load on the user when reading and interpreting the filter information.

The RD has been recently deployed in three different environments, one focusing on content consumption scenarios (Browser Extension) and two focusing on content creation tasks , such as writing paper or reports. We are currently collecting usage logs which will be evaluated to understand which features are actually being used, how the typical workflows look like, and whether they differ significantly from what we have envisioned. To evaluate the usefulness of the micro-visualisations and understand how their performance and incured user load compares to classical text-based faceted interfaces, we are planing to conduct a user study in a controlled environment. In the study we will focus on collaborative bookmarking scenarios, such as students collecting related work on a particular research area. In this context we will examine the relevant properties of readability, which is important when handling tasks over to other users, and memorability, which is crucial when continuing a task started days earlier.

In the long term we plan add interactivity to the micro-visualisations and explore ways to realise a version of the RD optimised for small screen mobile devices.

Acknowledgements: This work was funded by the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement Nr 600601. The Know-Center GmbH is funded within the Austrian COMET Program - managed by the Austrian Research Promotion Agency (FFG).

References

[1] Lee, B. and Smith, G. and Robertson, G. G., and Czerwinski, M. and Tan. D. S., FacetLens: exposing trends and relationships to support sensemaking within faceted datasets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '09). ACM, New York, NY, USA, pp.1293–1302, 2009.

[2] Delone, W. H. and McLean, E. R, The DeLone and McLean model of information systems success: a ten-year update, Journal of management information systems, 14, pp.9–30, 2003.

[3] M. Dörk and S. Carpendale and C. Collins and C. Williamson, VisGets: Coordinated Visualizations for Web-based Information Exploration and Discovery, IEEE Transactions on Visualization and Computer Graphics, 14, pp.1205–1212, 2008.

[4] Fried, J. Y. and van Iersel, M. P. and Aladjem, M. I. and Kohn, K. W. and Luna, A. PathVisio-Faceted Search: an exploration tool for multi-dimensional navigation of large pathways, Bioinformatics, 2013.

[5] Veeraraghavan, H. and Miller, J. V., Faceted Visualization of Three Dimensional Neuroanatomy By Combining Ontology with Faceted Search, Neuroinformatics, pp.245–259 2014.

[6] English, J. and Hearst, M. and Sinha, R. and Swearingen, K. and Lee, KP., Flexible search and navigation using faceted metadata, Technical report, University of Berkeley, 2002.

[7] Kern, R. and Jack, K. and Granitzer, M., Recommending Scientific Literature: Comparing Use-Cases and Algorithms. CoRR abs/1409.1357. 2014.

[8] Kienreich, W. and Lex, E. and Seifert, C., APA Labs: An Experimental Web-Based Platform for the Retrieval and Analysis of News Articles. In *Proceedings of the first International Conference on the Applications and Digital Information and Web Technologies (ICADIWT08)*, 2008

[9] Kules, B and Capra, R. and Banta, M. and Sierra, T., What do exploratory searchers look at in a faceted search interface?. In *Proceedings of the 9th ACM/IEEE-CS joint conference on Digital libraries (JCDL '09)*. ACM, New York, NY, USA, 313-322. 2009.

[10] Lex, A. and Schulz, H. and Streit, M. and Partl, C. and Schmalstieg, D., VisBricks: Multiform Visualization of Large, Inhomogeneous Data, IEEE, Vol. 17, Number 12 pp.2291–2300, 2011.

[11] Rauch, M. and Klieber, W. and Wozelka, R. and Singh, S. and Sabol, V., Knowminer Search - A Multi-visualisation Collaborative Approach to Search Result Analysis. 19th International Conference on Information Visualisation, pp.379–385, 2015.

[12] Roberts, J. C., Multiple view and multiform visualization, Proc. SPIE, 176–185, 2000.

[13] Schlötterer, J. and Seifert, C. and Granitzer, M. Web-based Just-In-Time Retrieval for Cultural Content. In Conf.proc. PATCH'14, 2014.

[14] di Sciascio, C. and Sabol, V. and Veas, E., uRank: Exploring Document Recommendations through an Interactive User-Driven Approach, RecSys Joint Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS'15), 2015.

[15] Seifert, C. and Jurgovsky, J. and Granitzer, M. FacetScape: A Visualization for Exploring the Search Space, in Conf.proc. IV'14, IEEE, pp.94–101, 2014.

[16] Ulbrich, E. and Veas, E. and Singh, S. and Sabol, V., Reading through Graphics: Interactive Landscapes to Explore Dynamic Topic Spaces, in proceedings of the 17th International Conference on Human-Computer Interaction (HCI International 2015). 2015.

[17] Tschinkel, G. and di Sciascio, C. and Mutlu, B. and Sabol, V., The Recommendation Dashboard: A System to Visualise and Organise Recommendations. In *proceedings of IV 2015*, pp. 241–244, 2015.