Controlling Personalized Recommendations in Two Dimensions with a Carousel-Based Interface

Behnam Rahdari[1], Peter Brusilovsky[1] and Alireza Javadian Sabet[1]

[1]School of Computing and Information, University of Pittsburgh, Pittsburgh, USA

Abstract

Using carousels to present recommendation results has been widely adapted for consumer-focused applications such as recommending movies and music. Carousel-based interfaces engage users in the recommendation process, leaving it to the user to decide which category of items is most relevant to them, yet leaving it to AI to produce a ranking of both items and carousels. This paper explores the idea to give possible uses of a carousel interface for two dimensions of user control, engaging then into both item ranking and carousel selection. We present an implementation of this idea for a recommender system that assists college students in finding relevant courses and advisors.

Keywords

Exploratory Search, Information Exploration, Open User Model, Recommendation System, Intelligent User Interface, Knowledge Graph, Carousel

1. Introduction

In this paper we present *Grapevine 2D*, a recommender system that assists university students in finding advisors and courses that match their interests and needs. Grapevine 2D combines ideas from the areas of user modeling, recommender systems, and exploratory searching to bring together traditional profile-based recommendation and extended opportunities for users to control the generation and presentation of recommended items. User control is important for recommender systems because immediate user needs and interests can't be reliably modelled. By providing a space for the user to influence the generation and presentation of recommendations, user-controlled recommenders could improve both the quality of results and overall user satisfaction [1, 2]. Despite considerable research on the topic, user-controlled recommender interfaces tend to focus on a just one "dimension" of user control such as controlling the engagement of peers [3, 1] or strength of contributing sources [4, 5]. In this paper, we present a carousel-based recommendation interface that offers users two complementary ways to control the generation and presentation of recommendation creating, or what we call, two dimensions of control. One dimension of control allows users to influence how recommended items in each

IntRS'21: Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, September 25, 2021, Virtual Event
ber58@pitt.edu (B. Rahdari); peterb@pitt.edu (P. Brusilovsky); alj112@pitt.edu (A. Javadian Sabet)
https://sites.pitt.edu/~ber58/ (B. Rahdari); https://sites.pitt.edu/~peterb/ (P. Brusilovsky); https://arjsabet.com/ (A. Javadian Sabet)
0000-0001-6514-912X (B. Rahdari); 0000-0002-1902-1464 (P. Brusilovsky); 0000-0001-9459-2411 (A. Javadian Sabet)

carousel are ranked by employing an open user model interface. Another dimension allows the user to control which carousels are presented and in what order. In contrast to the traditional "pure-AI" approach that focuses on generating "perfect" ranking of items and carousels despite insufficient information about user interests and preferences, a user-controlled approach allows both the "artificial" and "human" intelligence to collaborate and achieve the best results.

2. Related Works

2.1. Controllability in Recommender Systems

User controllability has been recognized as a valuable component of advanced information access interfaces. This research started with an exploration of user-controllable ranking of search results [6, 7] and later generated a stream of work on user-controllable recommender systems [3, 1, 2]. In the context of search, PeopleExplorer [8] offered users an option to re-sort people search results based on multiple user-related factors, while uRank [9] allowed then to control the importance of query terms. In the context of recommendation, a stream of research has explored the use of sliders to control various components of user profiles [1, 10] and the relative strength of contributing sources [4, 5]. More recently, *The Browser* [11] has offered a controllable visual interface for recommending articles to editors and a curation service for interesting writing; [12] investigated the extent to which interface element design can contribute to understanding, reflection, and modification of the recommendation result; and [13] explored the difficulties in using an interactive environment when multiattribute utility theory (MAUT)is used as one of the simplest methods for recommender systems.

Despite promising results demonstrated by the user-controlled recommender systems reviewed above, the majority of these explored controllable interfaces are reasonably complex and expect users to be relatively well prepared. As a result, few of these ideas have been adopted by the industry to date. In contrast, the industry has extensively adopted a very simple user-controlled recommender interface that is based on multiple carousels. Carousel-based interfaces allow recommender systems to deal with insufficient data to model user interests and preferences by presenting multiple "best guess" lists of results (carousels) for different scenarios and leaving it to the user to focus on one or more carousels to select their desired items. Despite the popularity of carousel-based interfaces in industry, research on carousel-based recommender interfaces in academia is in its early stage [14, 15, 16].

2.2. Open User Model for Personalized Information Access

The idea of applying open user models to better support information exploration process was among the first attempts to add transparency and control to user-adaptive systems. *Open* user models allow users to examine and possibly change the content of user models applied to personalize their search, browsing, or recommendation tasks. Despite a recognized success of open user models in the field of personalized learning (where these models are usually called open learner models [17]), the first attempts to introduce open user models to the area of personalized information access in the form of a keyword vector profile were not successful [6, 7]. However, a switch to semantic-level models that represented user interests over semantic entities, such as domain concepts [18, 19, 20] or named entities [21], have allowed several research teams to develop highly personalized information access systems that use open user models.

One problem of existing personalized information access systems with open user models is their complexity. All these system are based on relatively complex visualizations that usually present many dozens of visual objects to their users and require relatively complex manipulations for profile tuning. For example, Introspective Views [18] and SciNet [19] visualize open user models as a circular "radar" where more relevant concept points are located closer to the center. To tune these models, the users are allowed to move these points either toward the center or away from it. Adaptive VIBE [22] presents terms and named entities forming the user model as points of interests (POI) in a complex, relevance-based visualization. The fine-tuning could be done by docking and undocking POI and moving it within the visualization space. These actions lead to the user-adapted and user-controlled visual presentation of search results. While all these interfaces demonstrated their efficiency with advanced users, such as graduate students looking for relevant research papers, it has not yet been demonstrated that these complex interfaces could support a more typical information exploration process in which users are relative novices.

Our own experience with open user models focused on undergraduate students has demonstrated that coarse-grained open models composed of 10-20 elements [23] are very efficient and that they could support relatively complex activities, such as content navigation and social comparison. However, we also have strong evidence that fine-grained open models that include over 100 visual elements might be too complex for this category of users to interpret appropriately [24].

3. System Design

This section presents the interface and the underlying personalization mechanism of an information exploration and recommendation system, Grapevine 2D (an expansion of the earlier system Grapevine [25]), which is designed to help undergraduate and graduate students with no prior research experience in finding advisors and courses – for an undergraduate research project or a graduate study – that match their research interests. This task is known to be one that students find challenging, since they are frequently unable to express their research interests in terms that their prospective advisors use to describe their own work in their home pages, academic papers, or course descriptions. To support this task, Grapevine 2D helps every student to gradually form a model of their interests by discovering new research topics and keywords that match their not yet fully-formulated interests. Following the nature of the exploratory search, it provides users with different opportunities to recognize (rather than recall) relevant research topics in different contexts [26, 27, 19]. This model is then used to generate a series of semantically coherent carousels that each represent one aspect of user preference. It also remains visible to the students who could fine-tune it in the context of their exploration. The system's intelligent user interface is driven by a knowledge graph; a tightly connected network of research topics, prospective advisors, departments, publications and courses.

In the following, we describe how the Grapevine 2D system provides two complementary

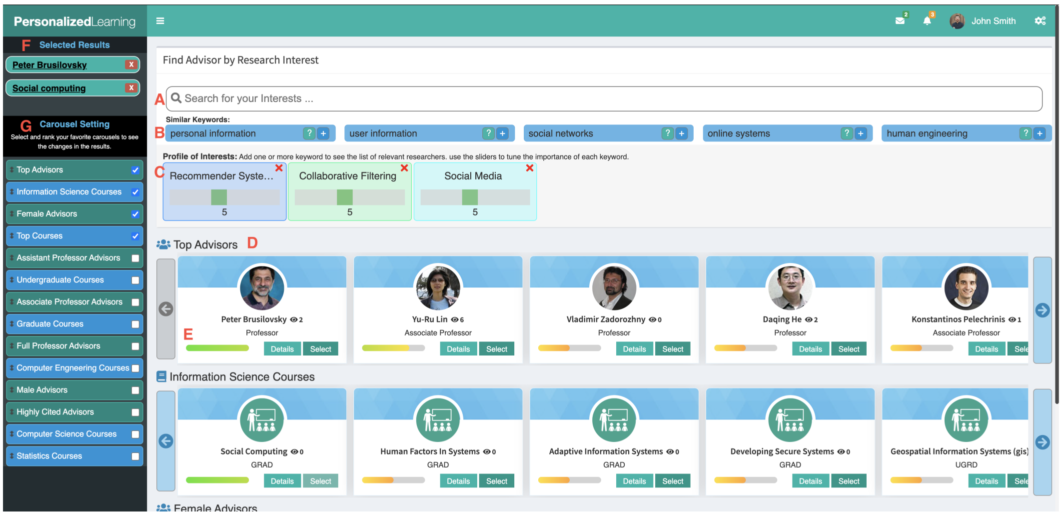

Figure 1: *Grapevine 2D* system interface. Features of the interface include: A: Search box; B: Recommended keywords; C: Interest profile and sliders; D: Results list; E: Heat bar; F: Final list; and G Carousel settings

ways to control the generation and presentation of recommendations to users.

3.1. Controlling Item Ranking Through an Open User Profile

The first dimension of interactivity centers around an "open user profile" that enables users to define and tune their interests.

The user profile of interest area (Figure 1C) is an open model of user interests. As a model of interests, it defines the system-generated list of recommended advisors and courses (Figure 1F); as an open model, it is visible and directly editable by the end users.

To edit the model, a user can add relevant topics, as well as remove less relevant keywords (using the red x), as they discover more relevant topics or explore different interests. Sliders associated with each keyword enable users to control the relative importance of a topic compared to others in their profile, ranging from 1 (least important) to 10 (most important). The use of sliders for the fine-tuning of a user profile was motivated by the keyword tuning approach in uRank design [9], which was confirmed as both user-friendly and efficient in a recommendation context. The initial value of the sliders is set to five, but can be changed at any time. All actions taken within the profile (adding, removing, or adjusting sliders) immediately affects the list of recommended items.

To initiate the recommendation process, users can use the search box (Figure 1A), which acts as the gateway to the system. Using an instant search approach, it allows users to discover relevant topics without a fully formulated query. When a user starts typing a query, a series of lexically similar keywords appears, which helps the user to discover a range of matching topics (e.g., Data Mining and Educational Data Mining). All keywords in the knowledge graph are considered when generating the list of instant search matches. To make finding a starting point easier for less prepared students, we also provided a series of typical topics of interest grouped in various semantic categories at the landing page of the Grapevine 2D system. Each key-phrase can be easily added to the user profile by clicking on the plus icon next to each item.

To further assist the user in the exploration process, when at least one keyword is added to the user's profile, a series of five semantically similar topics appear in the *Similar Keywords* area of the interface (Figure 1B). Users can add these recommended keywords to their interest profiles by clicking on the plus button to the right of each keyword. As the user's profile grows and is refined, the set of recommended keywords is updated, since the system recommends instances similar to all keywords in the user's profile. Each recommended keyword also provides users with a short description of the topic in question. Clicking on the question mark button next to the add button opens up a separate window containing the abstract of that keyword's Wikipedia entry. This information is crucial when the user is not familiar with the recommended keyword and needs more knowledge to decide whether the keyword should be added to the profile of interests.

Additionally, the Details view (Figure 2b) of recommended items provides users with additional information about items, such as an advisor's affiliation, research impact, research interests (shown as a list of topics), and external links to their research page and Google Scholar profile, as well as a description of the recommended courses. To stress an advisor's relevance to the user profile, the research interests that match the user profile are shown in green. Since faculty usually explore a set of related topics, we expect that the research interests of the prospective advisor that have not yet been added to the user profile (shown in blue) might also be relevant to the given user. To support the "interest discovery" process, these blue topic keywords can be added to the user's profile of interests and the short description of the topic can be reviewed with one click.

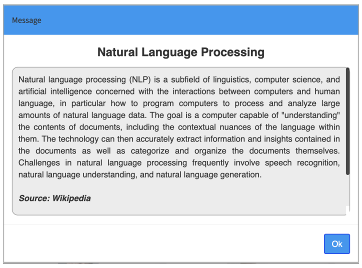

(a) Topic Explanation Window

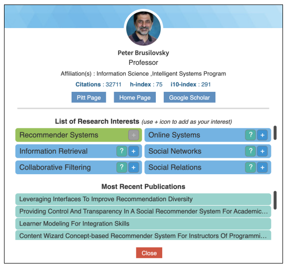

(b) Item Details View

Figure 2: (a) Topic explanation that contains the abstract of its Wikipedia entry, and (b) Detailed view showing advisor's academic profile

3.2. Controlling Preference Ranking Through a Carousel-Based Interface

The second dimension of controlability provides users with the ability to manually prioritize the more important aspects of recommended items. We achieve this by presenting the recommendation results in a series of semantically coherent carousels (Figure 1D) that each represent one aspect of advisors or courses.

Each carousel contains 10 items (advisors or courses) that are ranked by their relevance to the user's profile of interest and that share a particular aspect. (ex. Female Advisors or Graduate-Level Courses). Each Item in a carousel contains a photo and brief information about the advisor or course. There is also a relevance bar (Figure 1E) to reflect the item's relevance to the user profile of interests, along with two action buttons. The button *Details* opens the Details view (Figure 2b) that provides more details about the item. The button *Select* adds the item to the final list of results (Figure 1F). Throughout the search process, users can add multiple candidate items to the results list as they seem adequate to the user's needs at that point in time. Later, when the user's level of knowledge has changed by exploring the other options and learning about previously unknown topics, the results list will provide users with the opportunity to review their previous choices and check whether or not they still satisfy the new and updated criteria. The detailed view with information about the advisor is still accessible for advisors in the results by clicking on their names. If the advisor in the list is no longer meets the user's requirements, it can be simply removed by using the x button.

The carousel setting section on the left (Figure 1G) allow users to enable/disable or adjust the ranking of the carousel in the recommendation list. Users can enable/disable each carousel by using a checkbox on the right side of each item and can specify the order of a selected carousel through a simple drag-and-drop action. Any changes in the carousel setting will immediately affect the recommendation list saves for later visits. The ability to personalized the recommendation results based on important aspects of user interests allows users to focus their attention on exploring their interests and avoid dealing with items that are irrelevant to them (ex. a graduate student would prefer to see only graduate-level courses).

3.3. Knowledge Graph

Past experience with user-controlled personalized information access systems [28, 18, 22, 19] has demonstrated the importance of semantic entities to support a user's ability to understand and control personalization. These findings motivated us to choose a semantic-based knowledge graph as the core of the Grapevine 2D system personalized information access. The knowledge includes an expert curated list of available courses and advisors connected through a network of domain concepts (topics of interest) and hosted in a native graph database Neo4j. The knowledge graph is an underlying knowledge layer of the system that supplies the interface with real-time responses to user interactions.

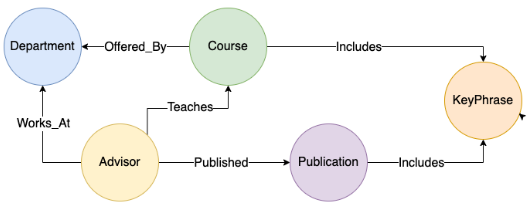

Figure 3: Graph schema that represents the entities of the knowledge graph and the relationship between them

To connect advisors with topics of interest, we extracted more than 4700 key phrases from nearly 7000 scientific articles authored by 180 researchers and 297 courses containing 472 key phrases in 27 schools and departments at the University of Pittsburgh. This collection consists of researchers and courses from a variety of disciplines related to information and computational science. The main target of this system is undergraduate and graduate students in the school of Computing and Information, but due to the diversity of degrees offered by the school, we included faculty and researchers from many other departments who could be relevant prospective advisors. Figure 3 presents the schematic representation of the knowledge graph.

To extract the key-phrases from the courses description and publications, we utilized the Computer Science Ontology (CSO) Classifier [29]; i.e., an ontology-driven topics detection for scholarly articles. The CSO [30] is a large-scale and granular ontology of the various computer science research areas. It is automatically generated using the Klink-2 algorithm [31] on a big dataset consists of approximately $16M$ articles, primarily in the field of computer science field. CSO encompasses more than $14K$ research topics and $162K$ semantic relationships among the topics. In addition to the main root, i.e., computer science, the ontology contains some secondary roots such as linguistics, geometry, and semantics. The CSO's automatic generation capability enabled us to easily extend the ontology, which is one of the important requirements for the Grapevine 2D system.

To extract the key phrases from the scholarly articles, we employ their *title* and *venue* to obtain their relevant key phrases. Similarly, to extract the key phrases from the courses, we feed the CSO Classifier with their *title* and *complete description.* The CSO Classifier is composed of two main components—namely, the *Syntactic module* and *Semantic module.* The Syntactic module seeks the CSO concepts that are *explicitly* mentioned in the input document, while the Semantic module at first employs part-of-speech tagging [32] to spot potential terms, then uses word2vec word embedding [33, 34] to build a list of inferred topics that are semantically related to the input document. In the next step, the resulting lists extracted by bth the Semantic and Syntactic modules will be merged together. Finally, the CSO classifier adds the relevant super-areas to the list.

3.4. Recommendation Method

We used the Cypher Querying Language to generate all the recommendations, including the main recommendations (advisors/courses) and relevant key phrase suggestions. For main recommendations that are presented in a carousel area, at each instance of user interaction with the interface (e.g., adding/removing keywords or changing the carousel settings), the system considers all advisors/courses connected to at least one of the topics of interest in the user profile and filters them based on the topic of each selected carousel. Then, a relevance score is assigned to each candidate that considers a candidate's similarity to each interest in the profile and the value of the sliders (Equation 1). Finally, the system ranks the candidates in their corresponding carousel by their relevance scores. If there are fewer than 10 candidates to recommend in each carousel, similar keywords are used to find more items.

In equation 1, A is a set of tuples $\{(a_1, w_1), (a_2, w_2), ...(a_n, w_n)\}$ that represent the current state of the user's profile (topics and weights) and f is a given item (advisor/course) in the graph. a_i and w_i correspond for i^{th} keyword and its slider value at the moment. $Sim_{(a_i, f)}$ shows the value

$$RelevanceScore_{(f,A)} = \sum_{i=0}^{size(A)} Sim_{(a_i,f)} * w_i \tag{1}$$

1: Calculation of relevance score for each candidate item

of relevance between a given keyword and a candidate item in our knowledge graph.

This value is pre-calculated using the cosine similarity between the keywords in an advisor's publications or a course's full description and its corresponding Wikipedia entry.

To generate the recommended keywords, the system generates three sets of candidate keywords for each set of keywords in the user's profile. These sets were created by using the co-occurrence of seed keywords and advisors' research interests, links, and categories (using collaborative filtering). Then, the system combines the number of co-occurred keywords in all three sets and uses it as a ranking mechanism. Finally, the system presents the top five results to the user.

4. Discussion and Future Works

In this paper, we presented a carousel-based interface that enables users to control personalized recommendations in two dimensions. The first dimension of interactivity allows users to define and tune the importance of their topics of interest using an open user profile that is used for generating recommendations. The second dimension of controllability provides control over the presented categories of recommended items by enabling users to select and rank relevant carousels.

Our next step is to conduct a formal user evaluation of the Grapevine 2D system to investigate the overall usability of using carousels for presenting and controlling the recommendation results. Furthermore, we are working on a new iteration of the Grapevine 2D system that employs a user interaction log to recommend relevant carousels and provide explanations for these recommendations using the knowledge graph. Last but not least, we intend to improve the knowledge graph by including more entities and relationships in order to produce better recommendations and to provide users with more carousel categories to choose from.

References

[1] B. P. Knijnenburg, S. Bostandjiev, J. O'Donovan, A. Kobsa, Inspectability and control in social recommenders, in: 6th ACM Conference on Recommender Systems, Association for Computing Machinery, New York, NY, USA, 2012, pp. 43–50.

[2] C. He, D. Parra, K. Verbert, 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.

[3] J. O'Donovan, B. Smyth, B. Gretarsson, S. Bostandjiev, T. Hšllerer, Peerchooser: visual interactive recommendation, in: Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems, ACM, 2008, pp. 1085–1088.

[4] D. Parra, P. Brusilovsky, User-controllable personalization: A case study with setfusion, International Journal of Human-Computer Studies 78 (2015) 43–67.
[5] C.-H. Tsai, P. Brusilovsky, Exploring social recommendations with visual diversity-promoting interfaces, ACM Transactions on Interactive Intelligent Systems 10 (2020) Article 5.
[6] A. Wærn, User involvement in automatic filtering: An experimental study, User Modeling and User-Adapted Interaction 14 (2004) 201–237.
[7] J.-w. Ahn, P. Brusilovsky, J. Grady, D. He, S. Y. Syn, Open user profiles for adaptive news systems: help or harm?, in: the 16th international conference on World Wide Web, WWW '07, Association for Computing Machinery, New York, NY, USA, 2007, pp. 11–20.
[8] S. Han, D. He, J. Jiang, Z. Yue, Supporting exploratory people search: a study of factor transparency and user control, in: Proceedings of the 22nd ACM international conference on Information & Knowledge Management, ACM, Association for Computing Machinery, New York, NY, USA, 2013, pp. 449–458.
[9] C. di Sciascio, V. Sabol, E. E. Veas, Rank as you go: User-driven exploration of search results, in: 21st International Conference on Intelligent User Interfaces, Association for Computing Machinery, New York, NY, USA, 2016, pp. 118–129.
[10] B. Rahdari, P. Brusilovsky, D. Babichenko, Personalizing information exploration with an open user model, in: Proceedings of the 31st ACM Conference on Hypertext and Social Media, HT '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 167–176.
[11] R. Bansal, J. Olmstead, U. Bram, R. Cottrell, G. Reder, J. Altosaar, Recommending interesting writing using a controllable, explanation-aware visual interface, in: Joint Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS'20) at 2020 ACM Conference on Recommender Systems, CEUR Workshop Proceeding, 2020, pp. 77–80.
[12] K. Herrmanny, S. Löppenberg, M. Schwarz, Investigating mechanisms for user integration in the activity goal recommendation process by interface design, in: Joint Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS'19) at 2019 ACM Conference on Recommender Systems, volume 2450 of *CEUR Workshop Proceeding*, 2019, pp. 46–54.
[13] Y. Nakao, T. Ohwa, K. Ohori, Generation of hints to overcome difficulty in operating interactive recommender systems, in: Joint Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS'19) at 2019 ACM Conference on Recommender Systems, volume 2450 of *CEUR Workshop Proceeding*, 2019, pp. 36–45.
[14] N. Felicioni, M. F. Dacrema, P. Cremonesi, A methodology for the offline evaluation of recommender systems in a user interface with multiple carousels, in: Adjunct Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization, 2021, pp. 10–15.
[15] D. Jannach, M. Jesse, M. Jugovac, C. Trattner, Exploring multi-list user interfaces for similar-item recommendations, in: Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization, ACM, 2021, pp. 224–228.
[16] L. Wu, M. Grbovic, J. Li, Toward user engagement optimization in 2d presentation, in: Proceedings of the 14th ACM International Conference on Web Search and Data Mining, ACM, 2021, pp. 1047–1055. doi:10.1145/3437963.3441749.

[17] S. Bull, Supporting learning with open learner models, in: 4th Hellenic Conference on Information and Communication Technologies in Education, volume A, Hellenic Conference in Information and Communication Technologies in Education, Athens, 2004, pp. 47–61.

[18] F. Bakalov, B. König-Ries, A. Nauerz, M. Welsch, Introspectiveviews: An interface for scrutinizing semantic user models, in: 18th International Conference on User Modeling, Adaptation, and Personalization (UMAP 2010), Springer, Big Island of Hawaii,USA, 2010, pp. 219–230.

[19] D. Glowacka, T. Ruotsalo, K. Konuyshkova, K. Athukorala, S. Kaski, G. Jacucci, Directing exploratory search: Reinforcement learning from user interactions with keywords, in: International conference on Intelligent User Interfaces, IUI '2013, Association for Computing Machinery, New York, NY, USA, 2013, pp. 117–127.

[20] B. Rahdari, P. Brusilovsky, Building a knowledge graph for recommending experts, in: Proceedings of the 1st International Workshop on Challenges and Experiences from Data Integration to Knowledge Graphs co-located with the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD 2019), Anchorage, Alaska, August 5, 2019, volume 2512 of *CEUR Workshop Proceedings*, 2019, p. 1.

[21] J.-w. Ahn, P. Brusilovsky, S. Han, Personalized search: Reconsidering the value of open user models, in: Proceedings of the 20th International Conference on Intelligent User Interfaces, Association for Computing Machinery, New York, NY, USA, 2015, pp. 202–212.

[22] J.-W. Ahn, P. Brusilovsky, Adaptive visualization for exploratory information retrieval, Information Processing and Management 49 (2013) 1139–1164.

[23] S. Sosnovsky, P. Brusilovsky, Evaluation of topic-based adaptation and student modeling in quizguide, User Modeling and User-Adapted Interaction 25 (2015) 371–424.

[24] S. M. Bailey, J. A. Wei, C. Wang, D. Parra, P. Brusilovsky, Cnvis: A web-based visual analytics tool for exploring conference navigator data, Electronic Imaging 2018 (2018) 376–1.

[25] B. Rahdari, P. Brusilovsky, D. Babichenko, Personalizing information exploration with an open user model, in: Proceedings of the 31st ACM Conference on Hypertext and Social Media, HT '20, New York, NY, USA, 2020, p. 167–176.

[26] M. A. Hearst, Clustering versus faceted categories for information exploration, Commun. ACM 49 (2006) 59–61.

[27] R. G. Capra, G. Marchionini, The relation browser tool for faceted exploratory search, in: Proceedings of the 8th ACM/IEEE-CS joint conference on Digital libraries, Association for Computing Machinery, New York, NY, USA, 2008, pp. 420–420.

[28] J.-w. Ahn, P. Brusilovsky, D. He, J. Grady, Q. Li, Personalized web exploration with task models, in: Proceedings of the 17th international conference on World Wide Web, 2008, pp. 1–10.

[29] A. A. Salatino, F. Osborne, T. Thanapalasingam, E. Motta, The cso classifier: Ontology-driven detection of research topics in scholarly articles, in: Digital Libraries for Open Knowledge, Springer International Publishing, Cham, 2019, pp. 296–311.

[30] A. A. Salatino, T. Thanapalasingam, A. Mannocci, F. Osborne, E. Motta, The computer science ontology: A large-scale taxonomy of research areas, in: The Semantic Web – ISWC 2018, Springer International Publishing, Cham, 2018, pp. 187–205.

[31] F. Osborne, E. Motta, Klink-2: Integrating multiple web sources to generate semantic topic

networks, in: M. Arenas, O. Corcho, E. Simperl, M. Strohmaier, M. d'Aquin, K. Srinivas, P. Groth, M. Dumontier, J. Heflin, K. Thirunarayan, K. Thirunarayan, S. Staab (Eds.), The Semantic Web - ISWC 2015, Springer International Publishing, Cham, 2015, pp. 408–424.

[32] A. Voutilainen, Part-of-speech tagging, The Oxford Handbook of Computational Linguistics (2012).

[33] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, arXiv preprint arXiv:1301.3781 (2013).

[34] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, in: C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, K. Q. Weinberger (Eds.), Advances in Neural Information Processing Systems, volume 26, Curran Associates, Inc., Red Hook, NY, 2013.