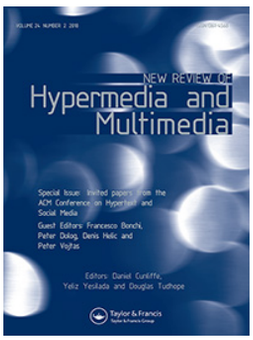

New Review of Hypermedia and Multimedia

ISSN: 1361-4568 (Print) 1740-7842 (Online) Journal homepage: http://www.tandfonline.com/loi/tham20

Invited papers from the ACM conference on hypertext and social media

Francesco Bonchi, Peter Dolog, Denis Helic & Peter Vojtas

To cite this article: Francesco Bonchi, Peter Dolog, Denis Helic & Peter Vojtas (2018) Invited papers from the ACM conference on hypertext and social media, New Review of Hypermedia and Multimedia, 24:2, 61-62, DOI: 10.1080/13614568.2018.1504520

To link to this article: https://doi.org/10.1080/13614568.2018.1504520

Published online: 16 Aug 2018.

NEW REVIEW OF HYPERMEDIA AND MULTIMEDIA
2018, VOL. 24, NO. 2, 61–62
https://doi.org/10.1080/13614568.2018.1504520

Taylor & Francis
Taylor & Francis Group

EDITORIAL

Invited papers from the ACM conference on hypertext and social media

Introduction to the special issue

This Special Issue presents three invited papers, selected from among the best contributions that were presented at the 2017 ACM International Conference on Hypertext and Social Media (HT 2017) held in Prague, Czech Republic on 4–7th July 2017. Since 1987, HT has successfully brought together leading researchers and developers from the Hypertext community. It is concerned with all aspects of modern hypertext research, including social media, adaptation, personalisation, recommendations, user modelling, linked data and semantic web, dynamic and computed hypertext, and its application in digital humanities, as well as with interplay between those aspects such as linking stories with data or linking people with resources.

The call for papers of HT 2017 was organised into four technical tracks: Social Networks and Digital Humanities (Linking people), Semantic Web and Linked Data (Linking data), Adaptive Hypertext and Recommendations (Linking resources), News and Storytelling (Linking stories). The Program Committee of HT 2017 accepted 19 papers (acceptance rate 27%) for regular presentation, and an additional 12 short-presentation papers. In addition, the conference featured four demonstrations and two keynotes: Kristina Lerman and Peter Mika.

The three papers selected for this Special Issue cover a diverse set of topics, well representing the spectrum of topics that were discussed at HT 2017.

The first paper, entitled “Implicit Negative Link Detection on Online Political Networks via Matrix Tri-Factorizations” (Ozer, Yildirim and Davulcu), deals with the prediction of negative connections between users of online political networks. Currently, the majority of social media sites do not support explicit negative links between participating users. However, the very nature of the political discourse often involves users in discussing controversial political issues, which results in a series of agreements and disagreements. The authors present a technically sound approach to extracting negative links from a variety of online political platforms by using a matrix factorisation approach. Matrix factorisation is extended in multiple ways to reflect the information that can be found in the sentiment of the written comments as well as the social balance theory known from the social sciences. The paper concludes with a range of experiments on the real datasets using the Twitter accounts of the politicians of the major UK political parties. The experiments show an improved accuracy of the community detection methods applied on the networks with the extracted negative interaction links as compared to the application of these methods on the networks having only positive links.

The second paper, entitled “Hybrid Recommendations by Content-Aligned Bayesian Personalized Ranking” (Peska) focuses on recommender systems that seek to predict the "rating" or "preference" a user would give to an item and hence enabling to display items in order the user might find interesting. A special problem is cold-start recommendation, i.e. for a new user or of a new item. The author proposes a hybrid recommendation technique “Content-Aligned Bayesian Personalized Ranking” (CABPR) with several variants. This is based on an existing Bayesian Personalized Ranking matrix factorization (BPR) by Rendle et al. CABPR

improves BPR's prediction by incorporating (arbitrary many) user-based or item-based similarity metrics. These similarities are utilized in the original BPR's optimization criterion, where additional regularization penalties are applied on similar users/items with dissimilar latent factors. The method also learns relevance of each type of similarity metric. An offline experimental evaluation on two datasets (movies, books) shows the effectiveness of the proposed approach. Datasets and software are available and so experiments are repeatable.

One of interesting future challenges mentioned by the author is the existence of a semantic gap between the external content-based attributes, such as director or actors for a particular movie (available in the majority of recommender systems datasets) and intrinsic (e.g. visual) features describing the movie itself. Combination of novel intrinsic objects' descriptors, e.g., originating from convolutional neural network, in case of multimedia content, with traditional collaborative filtering techniques may provide interesting results in the future.

The third paper, entitled "Characterizing Usage of Explicit Hate Expressions in Social Media" (Mondal, Silva, Correa and Benvenuto), describes research which should lead to better understanding of how hate speech can be detected and the reasons and conditions which lead to it in social media. The research aims at answering questions such as what is hate speech about, how does online and offline hate speech correlate, what role does anonymity play on hate speech, how do users perceive different categories of hate speech and how does hate speech vary across geography.

Regarding the first question, the paper makes use of sentence structure to detect hate speech. Using this approach, the article identifies number of hate intents (such as "I hate" or "I do not like") and number of hate targets (such as "Nigga", "White People" or "Negative People") together with their coverage in social media. The analysis is performed on Twitter and on Whisper. The article also categorises hate speech categories and their coverage. The top categories in both social media networks are "Race" and "Behavior". The article uses an FBI database to correlate between online and offline hate speech and finds a higher correlation between the database and Twitter and a lower correlation between the database and Whisper. The article also reports on the percentage of anonymous hate speech in each category. The paper utilises Anonymity Sensitivity Score to measure how users perceive hate speech. CrowdFlower is used to elicit annotations of the scores and Fleis kappa scores is used to check inter-annotator agreement concluding with slight to fair agreements. The paper also reports on differences between countries, for example, that hate speech on race in the US is higher than in Canada. The paper also reports on differences within the US. This paper is among the first to provide quantifications and insights on this matter.

We would like to thank the authors for their contributions to this Special Issue and the reviewers for their time and their insightful comments.

Francesco Bonchi
ISI Foundation, Turin, Italy
francesco.bonchi@isi.it

Peter Dolog
Department of Computer Science, Aalborg University, Aalborg, Denmark

Denis Helic
Institute for Knowledge Technologies, Graz University of Technology, Graz, Austria

Peter Vojtas
Department of Software Engineering, Charles University in Prague, Staré Město, Czechia