

Studying Confirmation Bias in Hashtag Usage on Twitter

Dominik Kowald
Know-Center GmbH
Graz, Austria
dkowald@know-center.at

Elisabeth Lex
Graz University of Technology & Know-Center GmbH
Graz, Austria
elisabeth.lex@tugraz.at

ABSTRACT

The micro-blogging platform Twitter allows its nearly 320 million monthly active users to build a network of follower connections to other Twitter users (i.e., *followees*) in order to subscribe to content posted by these users [9, 10]. With this feature, Twitter has become one of the most popular social networks on the Web and was also the first platform that offered the concept of *hashtags* (see post by Chris Messina in 2007¹). Hashtags are freely-chosen keywords, which start with the hash character “#”, to annotate, categorize and contextualize Twitter posts (i.e., tweets) [2, 13].

Problem. Although hashtags are widely accepted and used by the Twitter community, the heavy reuse of hashtags that are popular in the personal Twitter networks (i.e., own hashtags and hashtags used by followees) can lead to filter bubble effects [11] and thus, to situations, in which only content associated with these hashtags are presented to the user. These filter bubble effects are also highly associated with the concept of confirmation bias, which is the tendency to favor and reuse information that confirms personal preferences [12]. One example would be a Twitter user who is interested in political tweets of US president Donald Trump. Depending on the hashtags used, the user could either be stuck in a pro-Trump (e.g., “#MAGA”) or contra-Trump (e.g., “#fakepresident”) filter bubble.

Therefore, the goal of this paper is to study confirmation bias and filter bubble effects in hashtag usage on Twitter by treating the reuse of hashtags as a phenomenon that fosters confirmation bias.

Method. To reach this goal, we crawl two datasets from Twitter. The first one (i.e., *CompSci* dataset) consists of researchers from the field of computer science (= seed users) and their followees, while the second one (i.e., *Random* dataset) consists of random people (= seed users) and their followees. The statistics of these datasets are depicted in Table 1 and details of the crawling strategy are described in [8]. For all the seed users (i.e., $|U_S|$) in these datasets, we analyze (i) individual hashtag reuse (i.e., reusing own hashtags), and (ii) social hashtag reuse (i.e., reusing hashtags, which has been previously used by a followee) with respect to hashtag usage types and temporal effects.

Finally, we discuss ways to overcome confirmation bias and filter bubble effects in Twitter by means of recommendation mechanisms that focus on beyond-accuracy metrics.

Code. For reasons of reproducibility, we conduct this study using our open-source tag recommendation and benchmarking framework *TagRec* [4, 14], which is available via our GitHub repository².

Results. In our datasets, we first analyze hashtag assignments as well as hashtag reuse practices with the aim of studying confirmation

Dataset	$ U_S $	$ U $	$ T $	$ HT $	$ HTAS $
<i>CompSci</i>	2,551	91,776	5,649,359	1,081,403	9,161,842
<i>Random</i>	3,466	127,112	8,157,702	1,507,773	13,628,750

Table 1: Statistics of our *CompSci* and *Random* Twitter datasets. Here, $|U_S|$ is the number of seed users, $|U|$ is the number of distinct users, $|T|$ is the number of tweets, $|HT|$ is the number of distinct hashtags and $|HTAS|$ is the number of hashtag assignments. We aim to study confirmation bias in hashtag usage on Twitter by analyzing (i) the reuse rate of these hashtag assignments, and (ii) temporal effects of hashtag usage.




Figure 1: Analysis of hashtag reuse types as a proxy for studying confirmation bias. For each hashtag assignment, we study whether the corresponding hashtag has been used by the same user before in time (“individual”), by some of the followed users (“social”), by both (“individual/social”), by anyone else (“network”) or neither of them (“external”). We find that between 66% and 81% of hashtag assignments in both datasets can be explained by individual or social hashtag usage (i.e., the sum of “individual”, “social” and “individual/social”). This is a clear indication of confirmation bias in hashtag usage on Twitter.

bias in hashtag usage on Twitter. Specifically, for each hashtag assignment, we study whether the corresponding hashtag has either been used by the same user before (“individual”), by some of her followees (“social”), by both (“individual/social”), by anyone else in the dataset (“network”) or by neither of them (“external”). The results of this study are shown in Figure 1. We find that 66% of hashtag assignments in the *CompSci* dataset and 81% in the *Random* dataset can be explained by individual or social hashtag reuse, which is a clear indication of confirmation bias and filter bubble effects.

Next, we study to what extent temporal effects play a role in the reuse of individual and social hashtags on Twitter. Specifically, we analyze the recency of hashtag assignments (i.e., the time since the last hashtag usage). The results for this study are shown in the plots of Figure 2 and we find that temporal effects have an important

¹<https://twitter.com/chrismessina/status/223115412>

²<https://github.com/learning-layers/TagRec>

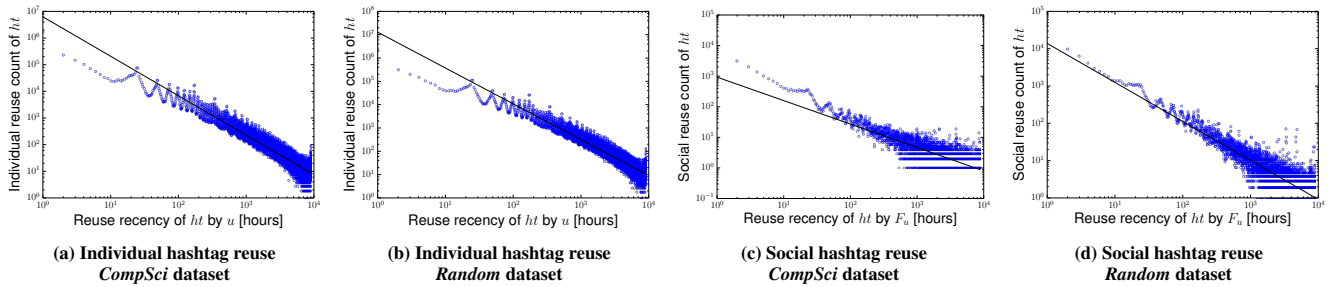


Figure 2: Temporal effects on individual and social hashtag reuse for the *CompSci* and *Random* datasets (plots are in log-log scale). Plots (a) and (b) show that the more recently a hashtag ht was used by a user u , the higher its individual reuse count (i.e., people tend to reuse hashtags that have been used very recently by their own). Plots (c) and (d) show that the more recently a user u was exposed to a hashtag ht , which was used by the followees F_u , the higher its social reuse count (i.e., people tend to reuse hashtags that have been used recently in the social network). We find that temporal effects play an important role in individual and social hashtag reuse, which is our measure for studying the confirmation bias in hashtag usage on Twitter.

influence on both individual as well as social hashtag reuse. People tend to reuse hashtags that were used very recently by their own and/or by their Twitter followees. Interestingly, there is a clear peak after 24 hours in all four plots, which further indicates that users typically (re)use the same set of hashtags in this time span and thus, tend to tweet about similar topics on a daily basis. This shows that also temporal effects are an important influence for confirmation bias and filtering bubble effects in hashtag usage on Twitter.

Conclusion & Discussion. Our findings of this paper are two-fold:

- (1) We find that between 66% and 81% of hashtag assignments in our datasets can be explained by individual or social hashtag reuse, which is a clear indication of confirmation bias in hashtag usage on Twitter.
- (2) We find that people tend to reuse hashtags that were used very recently by their own and/or by their Twitter followees, which means that confirmation bias in hashtag usage is also strongly influenced by temporal effects.

Based on this, we now discuss ways to overcome confirmation bias and filter bubble effects by using recommendation mechanisms. In our previous work (i.e., [8]), we presented a novel hashtag recommendation algorithm, which utilizes the Base-Level Learning (BLL) equation of the cognitive architecture ACT-R [1] for suggesting hashtags to users. Specifically, this algorithm ranks the hashtags of a user and her followees based on the factors of frequency and recency (see also [5, 7]). Although this approach achieves good results with respect to recommendation accuracy (see $BLL_{I,S}$ in Figure 3), it fosters hashtag reuse and thus, confirmation bias and filter bubbles effects in our two Twitter datasets.

In order to overcome these effects, we propose to focus on beyond-accuracy metrics of recommender systems such as diversity and serendipity [3, 6]. If we again take the example of a Twitter user who is interested in political tweets of US president Donald Trump, we could think of a hybrid recommendation approach with the goal to achieve a trade-off between accuracy, serendipity and diversity. For example, if we think of a pro-Trump Twitter user, a typical accuracy-focused recommender system would only recommend pro-Trump tweets and hashtags (e.g., “#MAGA”). Thus, our idea is

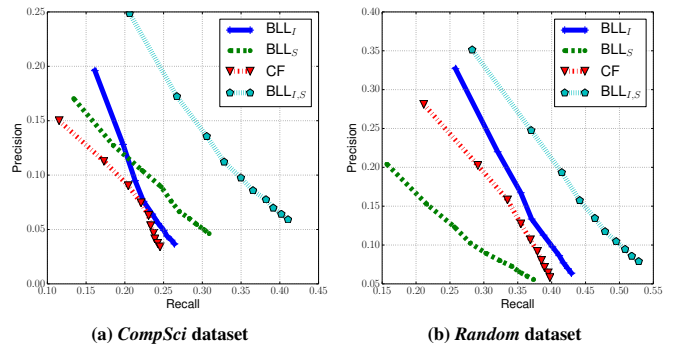


Figure 3: Precision / Recall plots showing the accuracy of BLL_I (i.e., only individual hashtags), BLL_S (i.e., only social hashtags), CF (i.e., Collaborative Filtering), $BLL_{I,S}$ (i.e., individual and social hashtags) for $k = 1 - 10$ recommended hashtags. $BLL_{I,S}$ provides the best results with respect to predication accuracy but also fosters the reuse of hashtags and thus, confirmation bias and filter bubbles effects in our two Twitter datasets.

to mix these pro-Trump recommendations also with contra-Trump recommendations (e.g., “#fakepresident”) to boost serendipity and diversity, and to give the user the chance to break the filter bubble.

Future Work. Finally, for future work, we also plan to verify our findings in larger Twitter data samples than the ones used in this paper as well as in other online social networks that feature the concept of hashtags, such as Instagram and Facebook.

Additionally, it would be interesting to evaluate the influence of social hashtags exposure, for example, by investigating retweet or mention networks in Twitter.

Acknowledgments. This work is funded by the Know-Center GmbH, which is part of the Austrian COMET program, and the H2020 project AFEL under grant agreement 687916.

Keywords. Twitter; Hashtag Usage; Confirmation Bias; Filter Bubble; Tag Recommendations; ACT-R; TagRec

REFERENCES

- [1] John R Anderson, Daniel Bothell, Michael D Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin. 2004. An integrated theory of the mind. *Psychological review* 111, 4 (2004), 1036.
- [2] Jeff Huang, Katherine M. Thornton, and Efthimis N. Efthimiadis. 2010. Conversational Tagging in Twitter. In *Proc. of HT '10*. ACM.
- [3] Marius Kaminskas and Derek Bridge. 2017. Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiIS)* 7, 1 (2017), 2.
- [4] Dominik Kowald, Simone Kopeinik, and Elisabeth Lex. 2017. The tagrec framework as a toolkit for the development of tag-based recommender systems. In *Proc. of UMAP'17*. ACM, 23–28.
- [5] Dominik Kowald, Simone Kopeinik, Paul Seitlinger, Tobias Ley, Dietrich Albert, and Christoph Trattner. 2015. Refining Frequency-Based Tag Reuse Predictions by Means of Time and Semantic Context. In *Mining, Modeling, and Recommending Things in Social Media*. Springer, 55–74.
- [6] Dominik Kowald and Elisabeth Lex. 2015. Evaluating Tag Recommender Algorithms in Real-World Folksonomies: A Comparative Study. In *Proc. of RecSys'15*. ACM, New York, NY, USA, 265–268.
- [7] Dominik Kowald and Elisabeth Lex. 2016. The Influence of Frequency, Recency and Semantic Context on the Reuse of Tags in Social Tagging Systems. In *Proc. of HT'16*. ACM, 237–242. DOI : <http://dx.doi.org/10.1145/2914586.2914617>
- [8] Dominik Kowald, Subhash Chandra Pujari, and Elisabeth Lex. 2017. Temporal effects on hashtag reuse in twitter: A cognitive-inspired hashtag recommendation approach. In *Proc. of WWW'17*. International World Wide Web Conferences Steering Committee, 1401–1410.
- [9] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. 2010. What is Twitter, a Social Network or a News Media?. In *Proc. of WWW '10*. ACM, New York, NY, USA, 591–600. DOI : <http://dx.doi.org/10.1145/1772690.1772751>
- [10] Seth A. Myers and Jure Leskovec. 2014. The Bursty Dynamics of the Twitter Information Network. In *Proc. of WWW'14*. ACM, New York, NY, USA, 913–924. DOI : <http://dx.doi.org/10.1145/2566486.2568043>
- [11] Eli Pariser. 2011. *The Filter Bubble: What the Internet Is Hiding from You*. (2011).
- [12] Scott Plous. 1993. *The psychology of judgment and decision making*. McGraw-Hill Book Company.
- [13] Daniel M. Romero, Brendan Meeder, and Jon Kleinberg. 2011. Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter. In *Proc. of WWW '11*. ACM, New York, NY, USA, 695–704. DOI : <http://dx.doi.org/10.1145/1963405.1963503>
- [14] Christoph Trattner, Dominik Kowald, and Emanuel Lacic. 2015. TagRec: towards a toolkit for reproducible evaluation and development of tag-based recommender algorithms. *ACM SIGWEB Newsletter Winter* (2015), 3.