

# The Influence of Frequency, Recency and Semantic Context on the Reuse of Tags in Social Tagging Systems

Dominik Kowald  
Know-Center  
Graz University of Technology  
Graz, Austria  
dkowald@know-center.at

Elisabeth Lex  
Knowledge Technologies Institute  
Graz University of Technology  
Graz, Austria  
elisabeth.lex@tugraz.at

## ABSTRACT

In this paper, we study factors that influence tag reuse behavior in social tagging systems. Our work is guided by the activation equation of the cognitive model ACT-R, which states that the usefulness of information in human memory depends on the three factors usage frequency, recency and semantic context. It is our aim to shed light on the influence of these factors on tag reuse. In our experiments, we utilize six datasets from the social tagging systems Flickr, CiteULike, BibSonomy, Delicious, LastFM and MovieLens, covering a range of various tagging settings. Our results confirm that frequency, recency and semantic context positively influence the reuse probability of tags. However, the extent to which each factor individually influences tag reuse strongly depends on the *type of folksonomy* present in a social tagging system. Our work can serve as guideline for researchers and developers of tag-based recommender systems when designing algorithms for social tagging environments.

## Keywords

social tagging; tag reuse; tag prediction; tag recommendation; frequency; recency; semantic context; ACT-R; BLL

## 1. INTRODUCTION

With the advent of the Web 2.0, social tagging has become an essential tool to collaboratively annotate content with freely chosen *tags* (i.e., keywords). The result of a social tagging process is a network connecting users, content resources and tags, which is referred to as *folksonomy* [9]. Since tags can be used to search for, browse and share content, they facilitate discovery and navigation in the Social Web [13]. Many social networks such as Medium, Twitter, Instagram and Facebook have adopted tagging in the form of hashtags as well [22].

Therefore, previous work has proposed tag recommendation algorithms with the aim to assist users in finding descriptive tags [4] and to control the shared tag vocabulary

[28]. In this respect, recent research has shown that a substantial amount of tag assignments can be explained by analyzing the information access in human memory [6, 24, 23], which is mainly influenced by past usage frequency, recency and the current semantic context [2, 1].

Even though there is already a large body of research available, which proposes algorithms for recommending tags [11, 20, 10, 18, 29, 21], none of these approaches incorporate all three aforementioned factors. For example, FolkRank [10] solely utilizes tag frequency and the current semantic context, whereas GIRP [29] solely builds on tag frequency and recency. Typically these methods have been evaluated as integrated models and thus, it remains unclear to what extent the individual factors of the approaches contribute to the final algorithmic performance.

### The present work: Factors that influence tag reuse.

In this work, we study factors that potentially influence the tag reuse behavior in social tagging systems (see also related work such as [8, 3, 7]). Specifically, we analyze the influence of frequency, recency and semantic context on the reuse of tags. Hence, it is our goal to better understand to what extent these factors can be exploited to predict the reuse of tags given a specific *folksonomy type*. This should lead to a guideline for designing and implementing tag prediction algorithms for given environments.

To that end, we integrate and extend our previous work on tag recommender systems [17, 14, 27], where we adapted the activation equation from the cognitive architecture ACT-R [1] to develop a model termed *BLLAC*. This model enables the prediction of future tag assignments for a user  $u$  by modeling the usefulness of a piece of information  $i$  – in our case, a tag – in  $u$ 's memory based on three factors: (i) how *frequent*  $i$  was used by  $u$  in the past, (ii) how *recent* (i.e., the time since the last usage)  $i$  was used by  $u$  in the past, and (iii) how useful  $i$  is for  $u$  in the current *semantic context*.

This is achieved by the two components of the activation equation: First, the base-level learning component (*BLL*), which integrates the factors of frequency and recency via a power function for reflecting the time-dependent decay of tag reuse [2] and second, the associated component (*AC*), which models the current semantic context as the similarity of tag  $i$  to tags already associated with the currently tagged resource  $r$ . However, since *BLLAC* utilizes the three factors, (i) frequency, (ii) recency, and (iii) semantic context, as an integrated model, it is still unclear to what extent these factors individually contribute to the efficacy of the model. Besides, we assume that the influence of these factors on predicting tag reuse depends on the folksonomy

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

HT '16, July 10–13, 2016, Halifax, Nova Scotia, Canada.

© 2016 ACM. ISBN 978-1-4503-4247-6/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2914586.2914617>

type, i.e., *narrow* (e.g., Flickr), *mixed*<sup>1</sup> (e.g., BibSonomy) or *broad* (e.g., MovieLens) [9], of the given social tagging system. This leads to the following two research questions:

- *RQ1*: How are the factors of frequency, recency and semantic context influencing a tag’s probability of being reused in social tagging systems?
- *RQ2*: Can the factors of frequency, recency and semantic context be exploited to efficiently predict a user’s tag reuse given a specific folksonomy type?

**Methods and findings.** In order to address *RQ1*, we conducted an empirical study on six social tagging datasets (i.e., Flickr, CiteULike, BibSonomy, Delicious, LastFM and MovieLens), in which we analyzed the influence of the three factors frequency, recency and semantic context on the reuse probability of tags (see Section 2). Next, to answer *RQ2*, we carried out a prediction study on the same datasets, in which we not only compared algorithms that reflect these factors individually but also approaches that combine these factors or incorporate social influences, e.g., by suggesting related tags of other users (see Section 3). We find that frequency, recency and semantic context positively influence the reuse probability of tags in all systems (*RQ1*) and that the efficacy of these factors for predicting a user’s tags depends on the folksonomy type of the given system (*RQ2*).

**Significance of this work.** To the best of our knowledge, this is the first study of its kind, which analyzes the reuse of tags to provide a transparent overview of the factors that influence the prediction of tags and, at the same time, relates these factors to the folksonomy type of the given system. Our findings may serve as a guideline for researchers and developers of tag-based recommender systems with regard to choosing the right prediction and recommendation methods for given social tagging environments.

## 2. EMPIRICAL STUDY

In this section, we present the datasets, methodology and results of our empirical study carried out to address *RQ1*.

### 2.1 Datasets

For the sake of this study, we turn to publicly available, real-world datasets gathered from the six social tagging systems Flickr<sup>2</sup>, CiteULike<sup>3</sup>, BibSonomy<sup>4</sup>, Delicious<sup>2</sup>, LastFM<sup>5</sup> and MovieLens<sup>5</sup>. To make our results comparable and to ensure reproducibility, we utilize the exact same dataset samples that were used in the study of [16]. A major advantage of these dataset samples is that they were created without any *p*-core pruning technique to ensure an unbiased evaluation [5]. Additionally, these datasets represent social tagging systems of various domains (i.e., images, Web links, references, music and movies), and differ in size and narrowness degree. The narrowness degree is defined as the average number of posts per resource, which allows to distinguish between narrow (Flickr), mixed (CiteULike, BibSonomy, De-

<sup>1</sup>With *mixed* folksonomies, we denote folksonomies that cannot strictly be categorized into the narrow or broad setting.

<sup>2</sup><https://www.uni-koblenz.de/FB4/Institutes/IFI/AGStaab/Research/DataSets/PINTSExperimentsDataSets>

<sup>3</sup><http://www.citeulike.org/faq/data.adp>

<sup>4</sup><http://www.kde.cs.uni-kassel.de/bibsonomy/dumps>

<sup>5</sup><http://grouplens.org/datasets/>

Dataset	$ U $	$ R $	$ T $	$ P $	$ P / R $
Flickr	9,590	856,755	125,119	856,755	1.000
CiteULike	18,474	811,175	273,883	900,794	1.110
BibSonomy	10,179	683,478	201,254	772,108	1.129
Delicious	15,980	963,741	184,012	1,447,267	1.501
LastFM	1,892	12,522	9,748	71,062	5.674
MovieLens	4,009	7,601	15,238	55,484	7.299

**Table 1: Statistics of our datasets. Here,  $|U|$  is the number of users,  $|R|$  is the number of resources,  $|T|$  is the number of distinct tags,  $|P|$  is the number of posts and  $|P|/|R|$  is the degree of narrowness.**

licious) and broad (LastFM, MovieLens) folksonomies [9]. The final statistics of our datasets are shown in Table 1.

### 2.2 Methodology

In order to analyze the tag reuse behavior of users, we split our datasets into training and test sets via an evaluation protocol, which persists the chronological order of the data. Thus, for each user  $u$ , we sorted his/her  $n$  posts by time, allocated the  $n^{\text{th}}$  (i.e., the most recent) post to the test set and the first  $n - 1$  posts to the training set [16, 11].

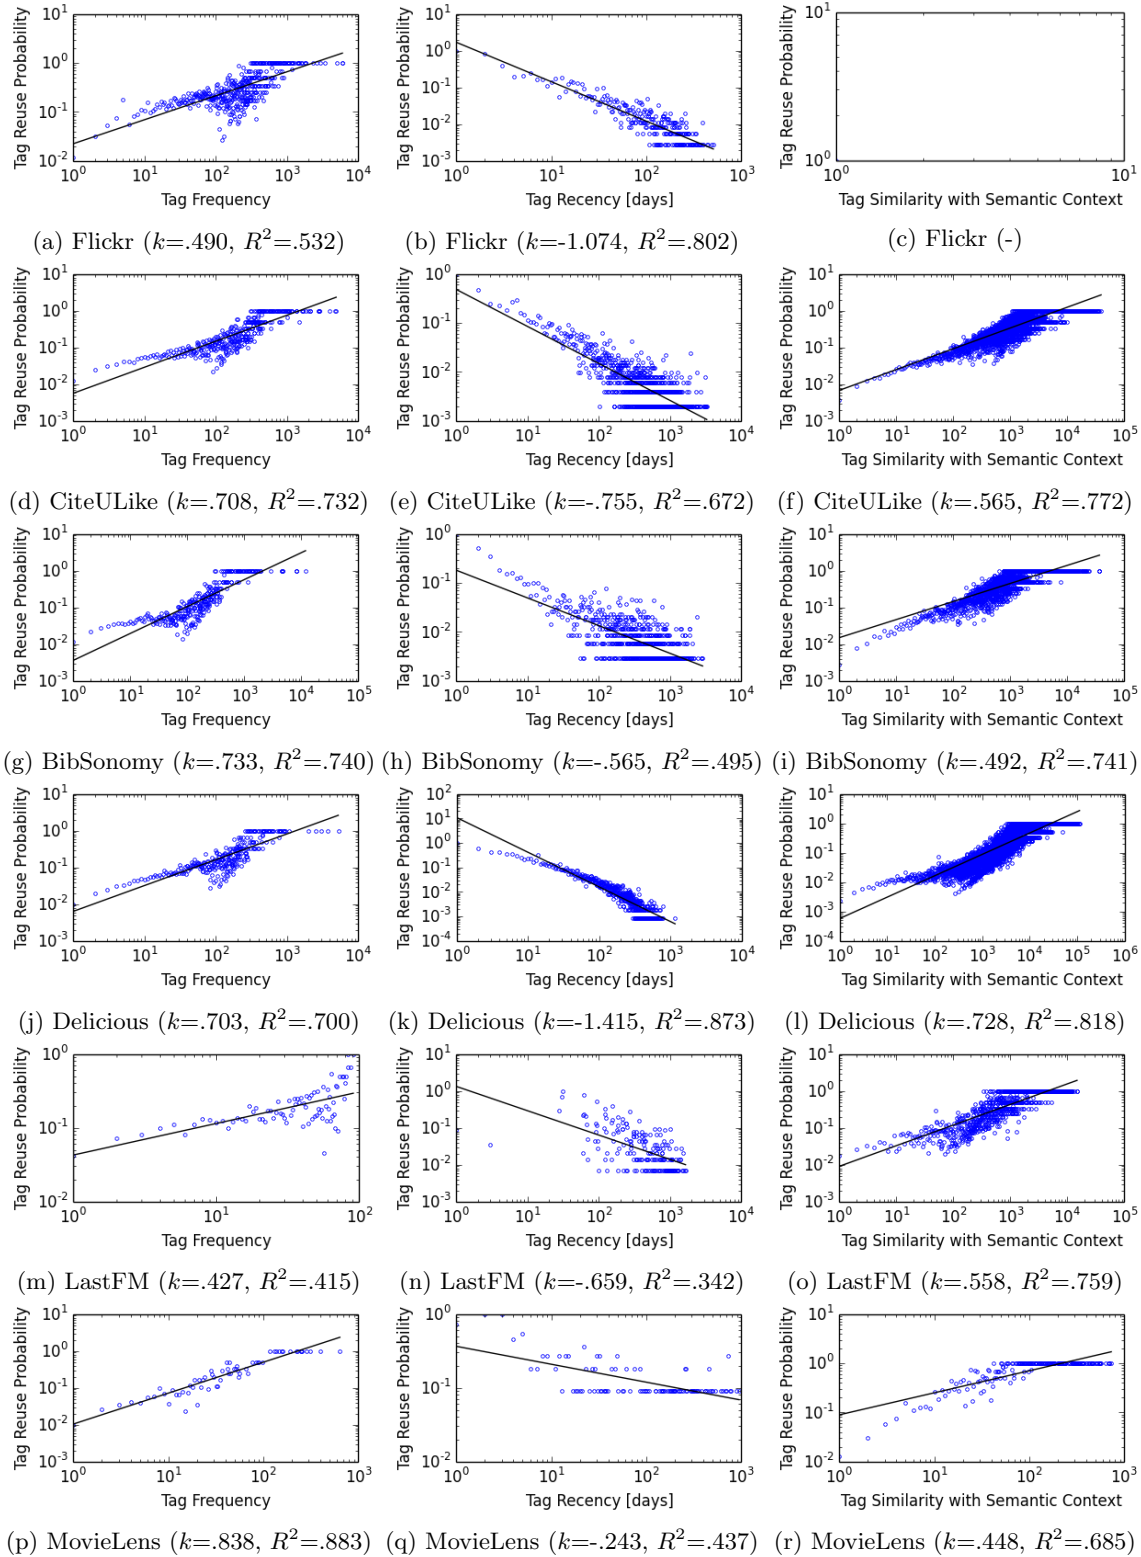
Next, in order to quantify the influence of frequency, recency and semantic context on the reuse of tags, we compared the tag assignments of the first  $n - 1$  posts in the training set (i.e., reflecting the past) of user  $u$  with the tag assignments of  $u$ ’s  $n^{\text{th}}$  post in the test set (i.e., reflecting the future). More specifically, for each tag  $i$  of  $u$ , we counted the number of times  $i$  was used by  $u$  in the training set (i.e., the frequency value) and determined if  $i$  was also reused by  $u$  in the test set (i.e., the reuse probability of  $i$  by  $u$ ). This procedure was repeated for the tags of each user  $u \in U$ . Finally, to obtain a statistically reliable value, we pooled together all tags with the same frequency value and calculated the mean reuse probability of these tags based on the individual ones.

We followed a similar procedure to study the influence of recency and semantic context. In the case of tag recency, we calculated the days elapsed since the last use of  $i$  by  $u$ . In the case of the semantic context, we determined the tag co-occurrence value between  $i$  and the tags already assigned to the currently tagged resource  $r$  (i.e., the second component of the activation equation). Then, as in the case of tag frequency, we pooled together the tags with the same recency or semantic context values and calculated the mean reuse probability based on the individual ones for both factors.

### 2.3 Results

In Figure 1, we plotted the mean tag reuse probability over (i) tag frequency, (ii) tag recency (in days), and (iii) tag similarity with the current semantic context on a log-log scale for our six datasets. We also provided the  $k$  (i.e., the slope) and  $R^2$  (i.e., the determination coefficient) estimates of the linear regression on the data. Specifically, we use the sign of  $k$  to determine how tag reuse is influenced by the factors and  $R^2$  to check if the data follows a power function. Across all six datasets, we can make three main observations:

1. The *more frequently* a tag was used in the past ( $k > 0$ ), the higher its reuse probability is.
2. The *more recently* a tag was used in the past ( $k < 0$ ), the higher its reuse probability is.
3. The *more similar* a tag is to tags in the *current semantic context* ( $k > 0$ ), the higher its reuse probability is.



**Figure 1: The influence of frequency, recency and semantic context on tag reuse in six social tagging systems (RQ1).** Clearly, all three factors positively influence a tag’s probability of being reused: Slope  $k > 0$  for frequency,  $k < 0$  for recency and  $k > 0$  for semantic context. Furthermore, the generally high  $R^2$  estimates of the linear regression indicate that large amounts of the data can be explained by a power function. *Please note* that there is no semantic context in Flickr since users solely tag their own resources in this system.

Folksonomy type	Dataset	Metric	Individual factors			Combination			Social	
			Frequency	Recency	SemCon	GIRP	BLL	BLL <sub>AC</sub>	FR	PITF
Narrow	Flickr	$F_1@5$	.371	.464	-	.455	<b>.470</b>	.470	.365	.350
		nDCG@10	.569	.702	-	.686	<b>.711</b>	.711	.561	.535
Mixed	CiteULike	$F_1@5$	.231	.236	.041	.243	.254	<b>.259</b>	.250	.178
		nDCG@10	.367	.385	.069	.394	.413	<b>.422</b>	.392	.294
	BibSonomy	$F_1@5$	.253	.252	.063	.262	.269	<b>.280</b>	.279	.215
		nDCG@10	.371	.368	.090	.386	.396	<b>.409</b>	.408	.327
Delicious	$F_1@5$	.173	.179	.108	.190	.203	<b>.243</b>	.196	.199	
	nDCG@10	.267	.287	.158	.298	.318	<b>.374</b>	.292	.302	
Broad	LastFM	$F_1@5$	.193	.189	.202	.198	.202	.251	.270	<b>.276</b>
		nDCG@10	.292	.293	.302	.303	.313	.375	.399	<b>.414</b>
	MovieLens	$F_1@5$	.077	.076	.077	.077	.079	.086	.153	<b>.156</b>
		nDCG@10	.177	.183	.176	.177	.187	.203	.319	<b>.324</b>

**Table 2: Prediction accuracy results of algorithms that (i) reflect the individual factors of frequency, recency and semantic context, (ii) combine these factors, and (iii) utilize social influences (RQ2). We see not only that all three individual factors can be exploited to efficiently predict a user’s tags but also that the efficacy of the algorithms depends on the folksonomy type (i.e., narrow, mixed, broad) of the given system.**

As such, all three factors (i.e., frequency, recency and semantic context) have a positive influence on a tag’s probability of being reused. Furthermore, the generally high  $R^2$  estimates in the log-log scaled plots indicate that large amounts of the data can be explained by a power function, as also indicated by the first component of the activation equation [2]. Thus, the first research question of our work (RQ1) can be answered affirmatively. In order to better understand the (individual) influence of each of the factors for predicting tags given a specific folksonomy type (RQ2), we conducted a prediction study described in the next section.

### 3. PREDICTION STUDY

In this section, we present the algorithms, methodology and results of our prediction study, which was conducted to address our second research question (RQ2). Specifically, we want to establish to what extent the factors of frequency, recency and semantic context can be exploited individually to efficiently predict a user’s tag reuse given a specific folksonomy type of a social tagging system.

#### 3.1 Algorithms

In terms of the compared algorithms, we utilize approaches that reflect the three factors individually, approaches that combine these factors and approaches that incorporate social influences (e.g., via related tags of other users).

**Individual factors.** To account for tag frequency, we utilize the Most Popular Tags ( $MP_u$ ) approach, which ranks the tags of a user based on their usage frequency [11]. To predict tags solely based on tag recency, we rank the tags of a user by the timestamp of their last usage. Since the third factor of interest, the semantic context  $SemCon$ , is represented by the second component of the activation equation, we recommend the tags that highly co-occurred with tags already assigned to the currently tagged resource [14].

**Combination of factors.** We utilize three algorithms based on combinations of the factors, which enables us to analyze if the performance of the individual factors can be improved when they are combined in the form of hybrid approaches. The first one, *GIRP* [29], integrates frequency and recency using an approach that models the effect of time (i.e., recency) via an exponential function. The second one, *BLL* [17], implements the first component of the activation equa-

tion and models the effect of time via a power function as suggested by [2]. The third algorithm in this respect, *BLL<sub>AC</sub>* [14], is the full implementation of the model, accounting for all three factors (frequency, recency and semantic context).

**Social influences.** To compare these methods that aim to predict a user’s individual tag reuse with approaches that also integrate social influences [25], we incorporate two well-known algorithms from tag recommender research in our study. The first one, *FolkRank (FR)* [10], is an extension of Google’s PageRank approach to iteratively rank the entities in folksonomies. The second one, *Pairwise Interaction Tensor Factorization (PITF)* [21], is based on factorization machines and has become one of the most successful methods for recommending tags. For both algorithms, we use the same parameter settings as in [16] (i.e., we set the preference vector weighting  $d$  for FR to .7 and the dimension of factorization  $\lambda$  for PITF to 256).

#### 3.2 Methodology

We conducted our prediction study on the same datasets and training/test set splits as in our empirical study (see Section 2.2). This allowed us to train the algorithms based on the posts in the training set and compare the top-10 tags that an algorithm predicted for user  $u$  and resource  $r$  of a post in the test set with the set of relevant tags in this test set post [11, 12]. Based on that, we computed two prediction accuracy metrics known from research on information retrieval and recommender systems. Specifically, we report the F1-score ( $F_1@5$ ) for the top-5 predicted tags<sup>6</sup> and the ranking-aware metric Normalized Discounted Cumulative Gain (nDCG@10) for the top-10 predicted tags [19].

To ensure the reproducibility of our results, we conducted this study via the open-source tag recommender evaluation and benchmarking framework *TagRec* [15, 26], which can be freely downloaded from GitHub<sup>7</sup> for scientific purposes.

#### 3.3 Results

In Table 2, we present the results of our prediction study indicated by the  $F_1@5$  and nDCG@10 metrics. Across all six datasets and both metrics, we observe three patterns of results based on the folksonomy type of the tagging systems.

<sup>6</sup> $F_1@5$  was also used as the main metric in PKDD’09: <http://www.kde.cs.uni-kassel.de/ws/dc09/evaluation>

<sup>7</sup><https://github.com/learning-layers/TagRec/>

**Narrow.** In the narrow folksonomy Flickr, the semantic context has no influence since users are solely tagging their own images. The results show that frequency and especially recency can be exploited to efficiently predict a user’s tag reuse in this narrow setting. Furthermore, these factors even outperform FR and PITF, the two algorithms that also utilize social influences by means of other users’ tags. When combining frequency and recency, we see that the accuracy of the strong recency factor can only be slightly improved in the case of BLL, which models the time component via a power function, and even decreases in the case of GIRP, which builds on an exponential temporal decay function.

**Mixed.** For the mixed folksonomies CiteULike, BibSonomy and Delicious, we observe a good performance for the factors of frequency and recency, and an average one for the semantic context. Additionally, the results suggest that a combination of all three factors in the form of  $BLL_{AC}$  provides the highest accuracy estimates and outperforms FR and PITF. Again, BLL (and thus, the power function) is apparently better suited to combine frequency and recency than GIRP (and thus, the exponential function).

**Broad.** Interestingly, the algorithms in the broad folksonomies LastFM and MovieLens had a completely different behavior. In these datasets, since there are a lot of tags assigned by other users to the currently tagged resource, the semantic context is a much more important factor for predicting tags than in the narrow and mixed settings. Similarly to the narrow case, the combination of the factors only slightly improves the accuracy of the individual factors. Due to a high number of average posts per resource (5.674 for LastFM and 7.299 for MovieLens – see Table 1), FR and PITF, which utilize related tags of other users as well, provide the best results in the broad setting.

**Summary.** Our results are summarized in Table 3. Overall, they confirm that the factors of frequency, recency and semantic context can be exploited to efficiently predict a user’s tag reuse. Which factor to consider, however, strongly depends on the folksonomy type of the given social tagging system. As such, the second research question of our work (RQ2) can also be answered affirmatively.

#### 4. CONCLUSION AND FUTURE WORK

In this paper, we analyzed the influence of frequency, recency and semantic context on the reuse of tags in social tagging systems. To that end, we conducted an empirical study and a prediction study on datasets gathered from the six social tagging systems Flickr, CiteULike, BibSonomy, Delicious, LastFM and MovieLens. The empirical study aimed to answer our first research question of this work (RQ1) and determine how the factors of frequency, recency and semantic context influence a tag’s probability of being reused. The three main findings of this analysis are: (i) the *more frequently* a tag was used in the past, the higher its reuse probability is, (ii) the *more recently* a tag was used in the past, the higher its reuse probability is, and (iii) the *more similar* a tag is to tags in the *current semantic context*, the higher its reuse probability is. This confirms that all three factors (i.e., frequency, recency and semantic context) have a positive influence on a tag’s probability of being reused.

Our prediction study was designed to determine to what extent these three factors can be exploited to efficiently

Folksonomy type	Frequency	Recency	SemCon	Comb	Social
Narrow	+/-	+	-	+/-	-
Mixed	+	+	+/-	+	+/-
Broad	+/-	+/-	+	+/-	+

**Table 3: Summary of our prediction accuracy results showing the performance of the algorithms based on the given folksonomy type. Please note that “+” indicates a good performance, “+/-” indicates an average performance and “-” indicates a poor performance of a factor/an approach in a specific setting.**

predict a user’s tag reuse given a specific folksonomy type (RQ2). With that regard, we applied not only algorithms that reflect these three factors individually (*Frequency*, *Recency* and *SemCon*) but also algorithms that combine these factors (*GIRP*, *BLL* and *BLL<sub>AC</sub>*) and incorporate social influences, i.e., related tags of other users (*FR* and *PITF*). We observed three patterns of results based on the folksonomy type of the datasets, which are summarized in Table 3. In the narrow case, the factor of recency is the most important one, whereas in the mixed case all three factors highly contribute to the prediction accuracy and thus, the combination of these factors (*Comb*) provides the best results. Finally, in case of the broad folksonomies, the semantic context becomes a much more important factor than in the other two settings. Furthermore, the best results in this setting are obtained for PITF, which utilizes social influences (*Social*), such as related tags of other users, as well.

Overall, our results show that frequency, recency and semantic context positively influence the reuse probability of tags in all six datasets (RQ1) and that the efficacy of these factors in terms of predicting a user’s tags depends on the folksonomy type of the given social tagging system (RQ2). Additionally, with our findings summarized in Table 3, we provide a transparent overview of the factors that influence the prediction of tags based on the given folksonomy type. We believe that our findings may be a guideline for researchers and developers in the area of tag-based recommender systems, helping them to choose the correct prediction method for a given social tagging environment.

**Limitations and future work.** One limitation of this study is that it is mainly focused on the factors that influence the *individual* reuse of tags. Thus, for future work, we would like to extend this by also analyzing the influence of the reuse (i.e., imitation) of other users’ tags [28] and study how this *social influence* effects tag predictions. Furthermore, we plan to expand our model of the *semantic context*, since to date, we have analyzed it solely based on the tags that were already assigned to the currently tagged resource. In this respect, we want to incorporate other resource-dependent information, such as the resource title or content, as well.

Finally, we plan to enhance our guideline with results for other types of social systems that utilize the concept of tags. To that end, we would like to determine the influence of frequency, recency and semantic context in systems that incorporate *hashtags* (e.g., Twitter, Facebook and Instagram). At the same time, this would allow us to verify that our findings can be generalized to various types of tags.

**Acknowledgments.** The authors would like to thank Paul Seitlinger, Tobias Ley, Ilire Hasani-Mavriqi and Emanuel Lacic for valuable comments on this paper. This work is funded by the Know-Center and the EU-IP Learning Layers.

## 5. REFERENCES

- [1] J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin. An integrated theory of the mind. *Psychological review*, 111(4):1036, 2004.
- [2] J. R. Anderson and L. J. Schooler. Reflections of the environment in memory. *Psychological science*, 2(6):396–408, 1991.
- [3] C. Cattuto, V. Loreto, and L. Pietronero. Semiotic dynamics and collaborative tagging. *Proceedings of the National Academy of Sciences*, 104(5):1461–1464, 2007.
- [4] K. Dellschaft and S. Staab. Measuring the influence of tag recommenders on the indexing quality in tagging systems. In *Proc. of HT'12*, pages 73–82. ACM, 2012.
- [5] S. Doerfel and R. Jäschke. An analysis of tag-recommender evaluation procedures. In *Proc. of RecSys'13*, pages 343–346. ACM, 2013.
- [6] W.-T. Fu. The microstructures of social tagging: a rational model. In *Proc. of CSCW'08*, pages 229–238. ACM, 2008.
- [7] W.-T. Fu, T. Kannampallil, R. Kang, and J. He. Semantic imitation in social tagging. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 17(3):12, 2010.
- [8] S. A. Golder and B. A. Huberman. Usage patterns of collaborative tagging systems. *Journal of information science*, 32(2):198–208, 2006.
- [9] D. Helic, C. Körner, M. Granitzer, M. Strohmaier, and C. Trattner. Navigational efficiency of broad vs. narrow folksonomies. In *Proc. of HT'12*, pages 63–72. ACM, 2012.
- [10] A. Hotho, R. Jäschke, C. Schmitz, G. Stumme, and K.-D. Althoff. FolkRank: A ranking algorithm for folksonomies. In *LWA*, volume 1, pages 111–114, 2006.
- [11] R. Jäschke, L. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme. Tag recommendations in folksonomies. In *Knowledge Discovery in Databases: PKDD 2007*, pages 506–514. Springer, 2007.
- [12] R. Jäschke, L. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme. Tag recommendations in social bookmarking systems. *Ai Communications*, 21(4):231–247, 2008.
- [13] C. Körner, D. Benz, A. Hotho, M. Strohmaier, and G. Stumme. Stop thinking, start tagging: tag semantics emerge from collaborative verbosity. In *Proc. of WWW'10*, pages 521–530. ACM, 2010.
- [14] D. Kowald, S. Kopeinik, P. Seitlinger, T. Ley, D. Albert, and C. Trattner. Refining frequency-based tag reuse predictions by means of time and semantic context. In *Mining, Modeling, and Recommending 'Things' in Social Media*, pages 55–74. Springer, 2015.
- [15] D. Kowald, E. Lacic, and C. Trattner. Tagrec: towards a standardized tag recommender benchmarking framework. In *Proc. of HT'14*, pages 305–307. ACM, 2014.
- [16] D. Kowald and E. Lex. Evaluating tag recommender algorithms in real-world folksonomies: A comparative study. In *Proc. of RecSys'15, RecSys '15*, pages 265–268, New York, NY, USA, 2015. ACM.
- [17] D. Kowald, P. Seitlinger, C. Trattner, and T. Ley. Long time no see: The probability of reusing tags as a function of frequency and recency. In *Proc. of WWW'14 companion*, pages 463–468. International World Wide Web Conferences Steering Committee, 2014.
- [18] R. Krestel, P. Fankhauser, and W. Nejdl. Latent dirichlet allocation for tag recommendation. In *Proc. of RecSys'09*, pages 61–68. ACM, 2009.
- [19] M. Lipczak. *Hybrid tag recommendation in collaborative tagging systems*. PhD thesis, Dalhousie University Halifax, 2012.
- [20] L. B. Marinho and L. Schmidt-Thieme. Collaborative tag recommendations. In *Data Analysis, Machine Learning and Applications*, pages 533–540. Springer, 2008.
- [21] S. Rendle and L. Schmidt-Thieme. Pairwise interaction tensor factorization for personalized tag recommendation. In *Proc. of WSDM'10*, pages 81–90. ACM, 2010.
- [22] D. M. Romero, B. Meeder, and J. Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proc. of WWW'11*, pages 695–704. ACM, 2011.
- [23] P. Seitlinger and T. Ley. Implicit imitation in social tagging: familiarity and semantic reconstruction. In *Proc. of CHI'12*, pages 1631–1640. ACM, 2012.
- [24] P. Seitlinger, T. Ley, and D. Albert. Verbatim and semantic imitation in indexing resources on the web: A fuzzy-trace account of social tagging. *Applied Cognitive Psychology*, 2014.
- [25] H. Steck. Item popularity and recommendation accuracy. In *Proc. of RecSys'11*, pages 125–132. ACM, 2011.
- [26] C. Trattner, D. Kowald, and E. Lacic. Tagrec: towards a toolkit for reproducible evaluation and development of tag-based recommender algorithms. *ACM SIGWEB Newsletter*, Winter, 2015.
- [27] C. Trattner, D. Kowald, P. Seitlinger, T. Ley, and S. Kopeinik. Modeling activation processes in human memory to predict the use of tags in social bookmarking systems. *The Journal of Web Science*, 2(1):1–18, 2016.
- [28] C. Wagner, P. Singer, M. Strohmaier, and B. A. Huberman. Semantic stability in social tagging streams. In *Proc. of WWW'14*, pages 735–746. ACM, 2014.
- [29] L. Zhang, J. Tang, and M. Zhang. Integrating temporal usage pattern into personalized tag prediction. In *Web Technologies and Applications*. Springer, 2012.