

Purpose Tagging: Capturing User Intent to Assist Goal-Oriented Social Search

Markus Strohmaier
Graz University of Technology and
Know-Center
Inffeldgasse 21a
8010 Graz, Austria
markus.strohmaier@tugraz.at

ABSTRACT

The terms that are used by users during tagging have been found to be different from the terms that are used when searching for resources, which represents a fundamental problem for search in tagging based systems. To address this problem, we propose *purpose tagging* as a novel kind of tagging that focuses on capturing aspects of *intent* rather than *content*. By capturing the different purposes a given resource can serve, purpose tags appear useful to mediate between the vocabulary of user intent on one hand, and the vocabulary of contents and tags provided by social software applications on the other. The paper at hand makes the following contributions: 1) It extends the set of known kinds of tags with a novel type and 2) it provides first empirical evidence of the principle feasibility of purpose tagging and its potential to facilitate goal-oriented social search in an exploratory case study.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval; G2.3 [Discrete Mathematics]: Applications.

General Terms

Experimentation, Human Factors

Keywords

User intent, tagging, social bookmarking, social search

1. INTRODUCTION

1.1 Motivation

A recent study conducted by Heymann et al [6] found that the terms used to craft search queries are usually different from the terms that are used to tag resources in social software. One reason for this is assumed to be a difference in usage [6][8]:

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SSM '08, October 30, 2008, Napa Valley, California, USA.
Copyright 2008 ACM 978-1-60558-258-0 /08/10...\$5.00.

While search queries tend to express user intent [3], tags are usually non-intentional - they tend to focus on different aspects of *content* and hardly express the different purposes (or goals) a given resource can serve [5]. In the context of computer human interaction, this problem is known as the “*gulf of execution*” [15], which describes the cognitive gap between a user’s goals and a system’s functionality or content. To give an example: A user having the goal to “find a physician in Seattle” would not be served well by the current set of del.icio.us tags provided for the resource www.yellowpages.com. Among the top 25 tags for this resource, not a single entry relates to medical information (the 5 most popular tags for this resource are *directory*, *yellowpages*, *reference*, *phone*, *search*), and in the entire set of tags for this website, “*physician*”, “*doctor*” or “*medical*” is not included. Yet, yellowpages.com can be assumed to represent a perfectly adequate means to help achieving the goal “find a physician”.

This represents the crux of search in current tag-based applications: In order to find a website like yellowpages.com, a user having the goal to “find a physician” would be required to cognitively translate her goal into lower level concepts, such as specific tags (“*phonebook*”) or classification terms. This introduces a cognitive burden and a corresponding cost that prevents users from effectively achieving their goals on the web.

The majority of tags currently used in social software applications focuses on describing content (“*What it is*”) [5],[22] rather than intent (“*What it can be used for*”). This makes capturing intent, i.e. the different purposes a given resource can serve, from existing tags particularly challenging. Given the current state of tagging systems, an interesting question to ask is: *How can we capture and represent the different goals a given resource can serve in social software?*

The set of possible goals a given resource can serve is usually not constrained by its content, but by the agents interacting with it. To get back to the previous example: The resource www.yellowpages.com might serve the goal of “*finding a physician*”, but might equally well help to achieve a range of other goals such as “*contacting an old friend*”, “*organizing a marketing campaign*” or “*find others who share the same family name*”. Because intent relates more closely to the agents interacting with a given resource, we can not turn to traditional tags or the contents of the website for information. One way of bridging this “*gulf of execution*” [15] would require a novel form of tags that more closely resemble the different goals a given resource can help to achieve. This paper explores such an

approach by introducing and studying the concept of *purpose tags* (as opposed to other kinds of tags such as content tags or sentiment tags [22]).

This paper is organized as follows: In Section 2, we review the state of the art of folksonomy research. In Section 3 we introduce purpose tagging and a novel research prototype as a variation of existing tagging systems. Section 4 describes our research setup and the set of research questions we are exploring in this work. In Section 5 we present results from a case study that investigates the feasibility and utility of purpose tagging in the context of search in social bookmarking applications. Section 6 discusses related research and Section 7 concludes our work. The overall contribution of this paper is the introduction and empirical exploration of a novel type of tagging to facilitate goal-oriented social search.

2. A MODEL OF FOLKSONOMIES

Tagging in social software systems is often characterized by a tripartite graph with hyperedges. The three disjoint, finite sets of such a graph correspond to 1) a set of persons or users $u \in U$ 2) a set of resources or objects $o \in O$ and 3) a set of annotations or tags $t \in T$ that are used by users U to annotate objects O . A very general model of folksonomies is defined by a set of annotations $F \subseteq U \times T \times O$ (cf. [9],[13],[6],[7]).

However, it is interesting to observe that this general model of folksonomies represents a significant simplification of social software applications that is evident in a vast number of variations, including applications with different types of users U , tags T , objects O and combinations thereof. Changing these parameters in a given social software application is likely to influence the kinds of tags produced from application usage. In order to better understand these effects, we discuss each of these parameters in some greater detail next.

2.1 Variability in the Set of Users U

Users play different roles in social software systems. Variability in the set of Users U represents the variety of users and roles who are involved in tagging objects in a particular system. While in some systems, the set of users U actively involved in tagging resources is (almost exclusively) composed of object creators (such as “photos” in flickr.com), in other systems the set of users actively tagging objects is largely composed of object consumers (such as “music” in last.fm) or some combination thereof (e.g. del.icio.us). Another source of variability lies in the distinction between broad and narrow folksonomies¹, capturing the balance between actively tagging creators and consumers. Further research suggests distinguishing between at least four user roles including 1) resource author, 2) resource collector 3) indexer or tagger and 4) searcher [21]. Beyond these simple distinctions, other user roles are conceivable. It can be hypothesized that the tags created by resource creators (for example, to promote his/her photos on flickr) differ significantly from the tags created by resource consumers (for example, to organize songs on last.fm) or the tags required by searchers. Such differences can be assumed to

have significant implications for studies of folksonomies and the utility of tags to improve search.

2.2 Variability in the Set of Tags T

The general model of folksonomies does not only provide room for varying types of users, but different types of tags as well. Recently, research has developed an interest in studying the different kinds of tags that can be found in social software applications. Golder and Huberman [5] for example identified six different functions of tags, including 1) identifying what a resource is about (e.g. “cars”), 2) Identifying what it is (“article”) 3) Identifying who owns it (“New York Times”) 4) Refining categories (“newspaper”) 5) Identifying qualities or characteristics (“funny”) 6) Self reference (“mystuff”) and 7) Task organizing (e.g. “TOREAD”). Different kinds of tags would have different implications on search (e.g. the tag “cars” vs. the tag “TOREAD”). Other research on categories of tags focuses on the motivation of users, and why they engage in tagging ([1], [12]). This includes categories such as 1) future retrieval 2) contribution and sharing 3) attract attention 4) play and competition 5) self presentation and 6) opinion expression [12]. Interestingly, most categories of tags focus on aspects related to content, and neither empirical studies [22] nor conceptual categories of tags identified by recent research ([1],[5]) include an explicit category of tags that allows to capture the different purposes a given resource can serve.

2.3 Variability in the Set of Objects O

An obvious variation in the tri-partite model of folksonomies takes place in O , the type of objects that are in the center of social software applications, also called “Objects of sociality”². These objects of sociality are the reason why people affiliate (or can be affiliated) with each other in social software applications. Objects O could take the form of songs (last.fm), pictures (flickr.com), URLs (del.icio.us), movies (youtube.com) publications (bibsonomy.org) or arbitrary other social objects such as events, trips, podcasts, places, etc. The effectiveness of tags to aid search in social software seems to depend on the choice of objects O in a given social software application. For example: whether the set of objects O is composed of textual or multimedia resources can be expected to influence the overall structure and nature of tags assigned in a given application, and thereby influence their utility for search.

2.4 Extending the Model of Folksonomies

While we can observe some degree of consensus regarding a general tri-partite model of folksonomies in the literature ([13],[9],[12]), even if different vocabulary and notations are used), variability in these models has hardly been considered or made explicit in the past. Recent research makes first steps towards addressing this variability ([5],[12]) on a conceptual level, yet there are no formalizations or comprehensive empirical studies to demonstrate differences or properties of variations in this model.

To inform our research, we propose to make variability in existing formalizations of folksonomies explicit via a simple extension of the general tri-partite model with a set of qualifiers

¹ http://personalinfocloud.com/2005/02/explaining_and_.html

² http://www.zengstrom.com/blog/2005/04/why_some_social.html

q, r, and s: $F \subseteq U_q \times T_r \times O_s$, where U_q qualifies the set of users U , T_r qualifies the set of tags T and O_s qualifies the set of objects O . This model enables us to characterize existing social software applications and introduces three basic parameters to influence resulting folksonomies.




Figure 1 Characterizing Different Social Software Applications with the Extended Model of Folksonomies

On the left hand in Figure 1, we can characterize the social software application 43things.com as a tri-partite graph that varies the general model of folksonomies with regard to the object (here: “goals”). In the middle, del.icio.us is characterized as a variation of folksonomies focusing on websites as the “objects of sociality”. The difference between broad and narrow folksonomies is illustrated on the right (flickr.com) by varying the type of users U and the type of objects O . On a general level, the extended model of folksonomies allows to vary a set of parameters and characterize different folksonomies along their basic dimensions.

In the following, we use the extended model to introduce a new category of tags (i.e. purpose tags) and a corresponding novel prototype of social software (intentional social bookmarking) that aims to capture the different purposes a resource can serve.

3. PURPOSE TAGGING

3.1 Purpose Tags

In order to capture purpose, we introduce purpose tags T_p based on the extended tri-partite model of folksonomies. In contrast to other kinds of tags [5] that aim to describe aspects of content, purpose tags aim to explicitly capture aspects of intent, i.e. the different contexts in which particular resources can be used. When assigning purpose tags, users are assumed to tag a website with a specific purpose or goal in mind. To give an example: When bookmarking the website www.facebook.com, a user might tag the resource with a purpose tag T_p “organize a high school reunion”. The intuition behind that is that purpose tags 1) expand the vocabulary of traditional tags and 2) help bridging the gulf between user intent expressed in search queries and the resources users expect to retrieve.

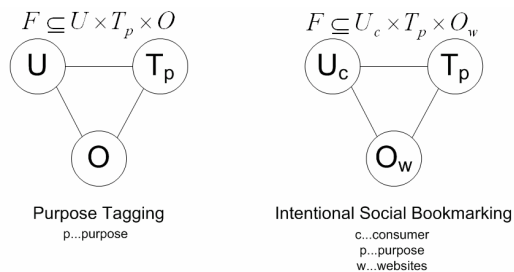


Figure 2 Purpose Tagging as a Variation of the Extended Model of Folksonomies

On a general level, the concept of purpose tagging does not make any assumptions about the kind of users U tagging objects O , nor the type of objects O that are being tagged. Using the extended model of folksonomies, we can characterize purpose tagging as a particular variation of the extended model in a way that is depicted on the left in Figure 2. To explore the feasibility and utility of this idea, we have developed a software prototype which is introduced next.

3.2 Prototype Implementation

We have developed a social bookmarking prototype with support for purpose tagging that represents a variation of “traditional” social bookmarking systems. The intentional social bookmarking prototype provides similar functionality that traditional social bookmarking systems usually provide (such as bookmarking websites, tagging them, constructing tag clouds, subscribing to user streams, etc), but focuses on gathering *purpose tags*. The system provides both a Firefox Plug-in and a so-called “Favelet” (Javascript code that is executed within a user’s browser) to support users in the bookmarking process. Purpose tagging is encouraged by adaptations to the user interface: When bookmarking a site, the interface requires users to complete the following sentence: “This resource helps me to: ...”. In contrast to traditional tagging user interfaces where users can assign arbitrary tags (such as “phone”, “book”, “directory”, etc), the prototype encourages users to formulate partial sentences that capture the different purposes of resources (such as “find a physician”) as opposed to their content.

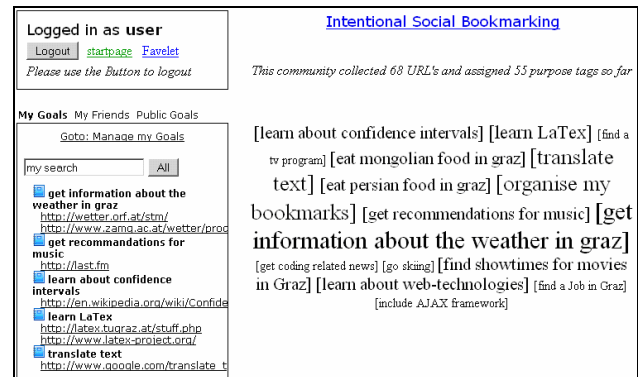


Figure 3 Screenshot of the Intentional Social Bookmarking Prototype

A “social” component allows users to add and copy URLs from other users (and corresponding purpose tags) to their own account. While our prototype is implemented as a variation of a social bookmarking application, there is nothing inherent in the concept of purpose tagging that would prevent it to be applicable to other types of applications as well (as depicted in Figure 2). Figure 3 shows a screenshot of the developed intentional social bookmarking system, depicting a purpose tag cloud on the right, and a list of purpose tags and resources of a given user on the left.

4. RESEARCH QUESTIONS AND SETUP

In an exploratory case study, we investigated purpose tagging by comparing data produced by two variations of social bookmarking applications: the purpose tagging prototype and del.icio.us, a popular social bookmarking system. We are

interested in exploring answers to the following research questions:

1. *Feasibility*: Would users assign meaningful purpose tags?
2. *Accuracy*: Do purpose tags accurately reflect plausible purposes of resources?
3. *Utility*: Can purpose tagging improve search in social software?
4. *Coverage*: Can purpose tags expand the vocabulary of existing tags?
5. *Meaning*: Are purpose tag graphs meaningful?

To explore answers to these questions, we have collected two different datasets (representing two variations on the general model of folksonomies F_A and F_B), which are introduced next.

4.1 Structure of Dataset A: $F_A \subseteq U_c \times T_p \times O_w$

Our first dataset was created by the purpose tagging prototype that requires users to assign purpose tags. In our folksonomy F_A , U_c refers to the *kinds of users* who are actively involved in bookmarking and tagging websites (*consumers*), T_p refers to *purpose tags* and O_w refers to *websites* that are bookmarked and tagged by users.

4.2 Structure of Dataset B: $F_B \subseteq U_c \times T \times O_w$

To compare this type of folksonomies with other kinds of tags assigned to resources, our second dataset was created by users of the social bookmarking system del.icio.us. The data was retrieved in a crawl of del.icio.us performed in July 2008. In F_B , U_c refers to the kinds of users who are actively involved in bookmarking and tagging websites (*consumers*), T refers to the entirety of tags assigned to O_w , and O_w refers to *websites* that are bookmarked and tagged by users.

Thus, the only structural variation between dataset A and B is the kind of tags that are used to annotate resources.

4.3 Study Design

Dataset A: For the case study, a set of 19 subjects (graduate students and employees of an applied research organization) was recruited to use the purpose tagging prototype over the course of two weeks, from June 4th to June 18th 2008, during their usual day-to-day online activities. The subjects were instructed to bookmark websites related to “Graz”, which is a small town with ~250.000 inhabitants located in Austria, Europe. We have chosen “Graz” because all subjects were familiar with the topic (all subjects study or work in Graz) and because it provides a sufficiently narrow topic to gather meaningful data during 2 weeks. Although none of the subjects had English as their first language, they were reasonably trained and were therefore instructed to assign English tags.

Dataset B: After acquiring dataset A, the set of URIs contained in Dataset A was used to identify a matching list of websites (based on exact URL matching) on del.icio.us. The resulting list of URIs was used to crawl all corresponding tags and users.

5. RESULTS

The case study yielded two datasets, which are depicted in Table 1. Over the course of two weeks, dataset A accumulated 123 URIs bookmarked by 19 users who assigned 75 purpose tags.

	URIs	Tags	Users
Dataset A (Purpose Tagging)	123	75	19
Dataset B (del.icio.us)	36	849	2801

Table 1 Characteristics of Dataset A and B

Out of the 123 URIs contained in dataset A, 36 URIs were bookmarked on del.icio.us as well. On del.icio.us, these 36 URIs were tagged by 849 unique tags and 2801 users. Based on these two datasets, we were interested in seeking answers to the five research questions introduced earlier.

5.1 Would Users Assign Meaningful Purpose Tags?

The case study (dataset A) produced a total of 75 unique purpose tags over two weeks. Figure 4 breaks this number down to individual users.

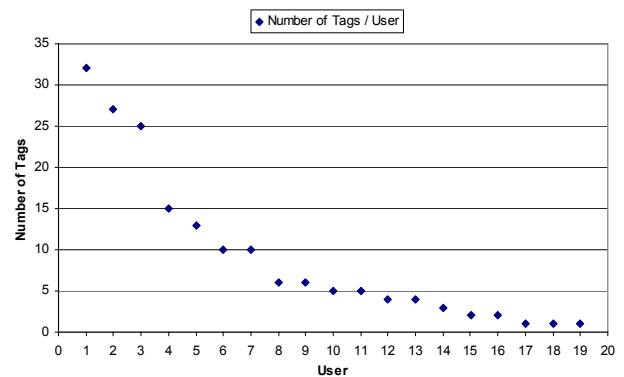


Figure 4 Activity of Users

The most active user in the dataset generated 32 purpose tags, the average amount of purpose tags assigned by a user is 9.05, the median is 5. Overall, users seemed to be comfortable assigning purpose tags to textual resources. To illustrate what kind of tags users have assigned, Figure 5 presents the top 20 tag-cloud produced by a simple weighted function. The weighted function considers the number of URLs tagged by a given purpose tag (20%) and the number of different users sharing this tag (80%) to scale the size of purpose tags. The larger the font size, the higher the score of the weighted function.

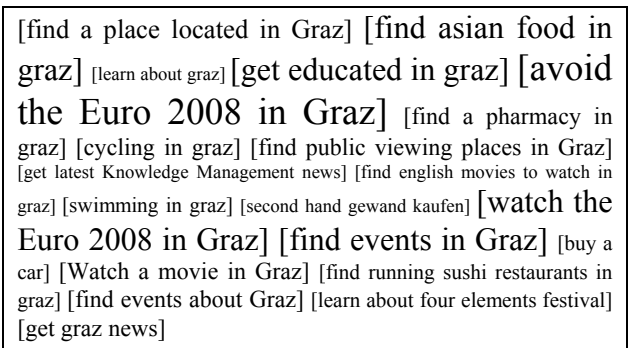


Figure 5 Purpose Tag Cloud Generated from Dataset A

Many of the purpose tags are related to a major event that took place in Austria in June 2008 – the European Soccer Cup EURO

2008 (“watch the EURO 2008 in Graz”, “find public viewing places in Graz”). The purpose tag that achieved the highest score in our dataset was “avoid the Euro 2008 in Graz”. This tag was assigned to a newspaper article for people who are not particularly interested in soccer. While the subjects of our study were instructed to exclusively assign tags related to Graz in English language, some used a mix of German and English (“second hand gewand kaufen”). Others seemed to use the prototype to collect resources not related to Graz as well (“get latest Knowledge Management news”). At a first glance, most purpose tags seem to represent plausible goals. But do the collected purpose tags accurately reflect the purposes of the URLs they are assigned to?

5.2 Do Purpose Tags Accurately Reflect Plausible Purposes of Resources?

To explore this question, we have conducted preliminary evaluations with 4 human subjects. The subjects did not use the purpose tagging prototype previously and were instructed to rate whether an arbitrary sample of 8 purpose tags drawn from the entire set accurately reflect a plausible purpose of the corresponding resources or not. The results are shown in Figure 6.




Figure 6 Perceived Accuracy of a Sample of 8 Purpose Tags

While there is some disagreement between the 4 human subjects, overall it seems that purpose tags retain their meaning across different user populations. The reasons for varying ratings provided included the level of detail or generality covered by the purpose tags. T3 (“learn about graz” for the resource www.graztourismus.at) and T5 (“use public transport” for the resource www.gvb.at) were rated inaccurate by two subjects. Although the 2 subjects expressed that these tags broadly reflect plausible purposes, they felt the tags could have been formulated in a more accurate way. All other tags received higher degrees of agreement.

5.3 Can Purpose Tagging Improve Search in Social Software?

Given the exploratory nature of our work on purpose tags, any ultimate judgement of their utility for improving search would be premature. However, in order to explore whether and how purpose tags have the potential to positively influence search in social software, we have conducted preliminary evaluations. The intentional social bookmarking prototype presented earlier implements rudimentary support for goal-oriented social search. The set of bookmarked websites can be searched via a simple substring matching of the search query with the set of purpose tags. If purpose tags would be able to better capture search

intent, this simple type of search should already provide advantages over search in traditional tagging systems. The search interface provides search results as users type their queries. Instead of listing all URLs related to the query, the search interface provides a list of purpose tags first, and users have to click on seemingly suitable tags in order to obtain the list of websites related to this purpose tag. This is depicted on the left hand side of Figure 7. The right hand side shows the traditional del.icio.us search interface, where 1) tags related to the query, 2) retrieved URLs and 3) the tags assigned to them, are listed.

Figure 7 Comparison of Search Results for Query “Graz”, left: Intentional Social Bookmarking, right: del.icio.us

Because the limited set of URIs available in both datasets (36) hinders meaningful quantitative evaluation, we have performed qualitative evaluations of the two different search interfaces by user observation. Four users were presented with four search tasks that included: “find an overview of restaurants in Graz”, “get a weather forecast for Graz”, “find information about local events in Graz”, and “find information about movie showtimes in Graz”. All search tasks (in principle) could be accomplished with the result sets generated by both search interfaces. During evaluation, interactions between the users and the search interfaces were captured via screen casts and audio recordings (with the exception of one user who refused recordings). The following observations were made:

While all users were comfortable using purpose tags to narrow down search results in the purpose tagging prototype, none of the users used del.icio.us tags (provided at the top of the search result page) to refine their search. Users reported that by being presented with a list of possibly related goals, they felt “guided” during search, and that purpose tags were helpful because they “felt natural” when trying to accomplish a search goal. All subjects noted that it was easier for them to assess whether a given purpose tag matches their search goal vs. assessing whether a given del.icio.us tag or a given URL from the result set matches it. This seems to hint towards the potential usefulness of purpose tags to guide and orient informational search. Indeed, during the search process, users used purpose

tags to refine and disambiguate their search intent, and to explore the set of associated URLs. One subject reported that he felt that some purpose tags are “misleading”. These were tags where the user did not agree on the relation between the purpose tag and the assigned resources.

Although the mechanism for search in the purpose tagging prototype was based on substring matching only, all users were successful in retrieving relevant purpose tags to refine their search goal. This was not always the case in the del.icio.us interface: Some users struggled with over-specified queries that yielded zero-size result sets. From this observation, it can be hypothesized that the vocabulary of purpose tags significantly expands the vocabulary of the kind of tags usually assigned on del.icio.us. This question is investigated next.

5.4 Can Purpose Tags Expand the Vocabulary of Existing Tags?

In order to explore this question, we have compared the set of tags assigned to 36 URIs that are contained in both datasets A and B with regard to their vocabulary size, that is the number of unique tokens contained in them. Considering that purpose tags consist of multiple tokens (e.g. “find a physician” contains 3 tokens), the vocabulary size of Dataset A is 92 (based on 38 purpose tags assigned to the 36 URIs). Assuming that each tag in the delicious tag set represents a single token, the vocabulary size of Dataset B is 849.

	URIs	Vocabulary Size	Users
(Subset) Dataset A (Purpose Tagging)	36	92	19
Dataset B (del.icio.us)	36	849	2801

Figure 8 Characteristics of Overlapping Subsets

Figure 8 gives an overview of the two datasets, and shows a proportional Venn Diagram illustrating the vocabulary overlap between the two tag sets.

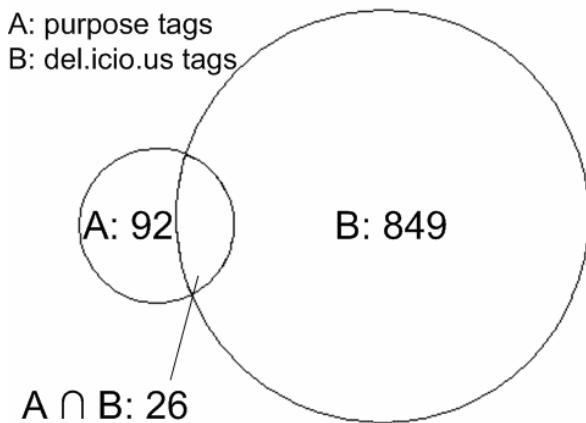


Figure 9 Vocabulary Overlap between Purpose Tags and Delicious Tags

The Venn diagram reveals that the majority of terms contained in the vocabulary of purpose tags ($66/92 = \sim 72\%$) is *not* contained in del.icio.us tags. To better understand the nature of

these sub sets, Table 2 shows the elements contained in $A \setminus B$ and $A \cap B$.

$A \setminus B$	$A \cap B$
pharmacy, to, via, location, buy, for, street, timetable, euro, where, asian, of, medical, more, results, facts, ticket, located, on, place, educated, transport., get, how, watch, offers, learn, practitioner, map., climb, specific, about, insight, use, from, running, swimming, figure, dates, detailed, tourists., relevant, restaurants, find, graz., airport., cityskating, b, news, a, at, into, the, learnpartner, in, know, dentist, flight, new, schedule, overview, what, arrival, public, first, out	car, 2008, knowledge, technology, cycling, events, visit, online, em2008, friends, transport, sushi, weather, train, conference, attention, metadata, information, university, movie, graz, tu, cinema, program, food, management

Table 2 $A \setminus B$ and $A \cap B$

The results suggest that purpose tags play a rather different role compared to other tags. While a few tags differ only in minor aspects (punctuation, e.g. “graz.”) or represent traditional stopwords (“a”, “at”, “into”), the majority of elements in $A \setminus B$ represent a true expansion of the vocabulary of B. This observation can also be corroborated by inspecting specific examples in the purpose tag set. In a pilot study, a user assigned the purpose tag “find a girlfriend” to the website www.studivz.net (a German-speaking version of facebook.com, not depicted in Table 2). While “find a girlfriend” seems a plausible goal that facebook.com can help to achieve, it would be hard to extract this kind of information from other kinds of tags, or the contents of the website itself. Among 21,104 users who were bookmarking facebook.com on del.icio.us, not a single person has used “girlfriend”, “girl-friend” or “girl_friend” as a tag. This anecdotal evidence combined with the evidence presented in Figure 9 seems to provide first evidence that purpose tags help to expand the vocabulary of existing kinds of tags in a useful way.

5.5 Are Purpose Tag Graphs Meaningful?

A common way to study folksonomies is to transform the tripartite graph U, T, O into bipartite graphs, which are also referred to as affiliation or two-mode networks. This allows constructing three different bipartite graphs: the UO , OT and UT graph. These graphs can be folded ([13],[9]) into six different unipartite graphs. Given a graph G representing the bipartite graph OT respectively UT , a graph T^* connecting tags based on co-occurrence with resources can be calculated via $T^* = G^T G$.

When studying semantics of tags, not all possible unipartite graphs are equally interesting, but two unipartite graphs are of special relevance: The two graphs T_1^* and T_2^* that can be folded from UT and OT allow to construct two different unipartite graphs where relations between tags are inferred from the tags’ co-occurrence with users U respectively objects O .

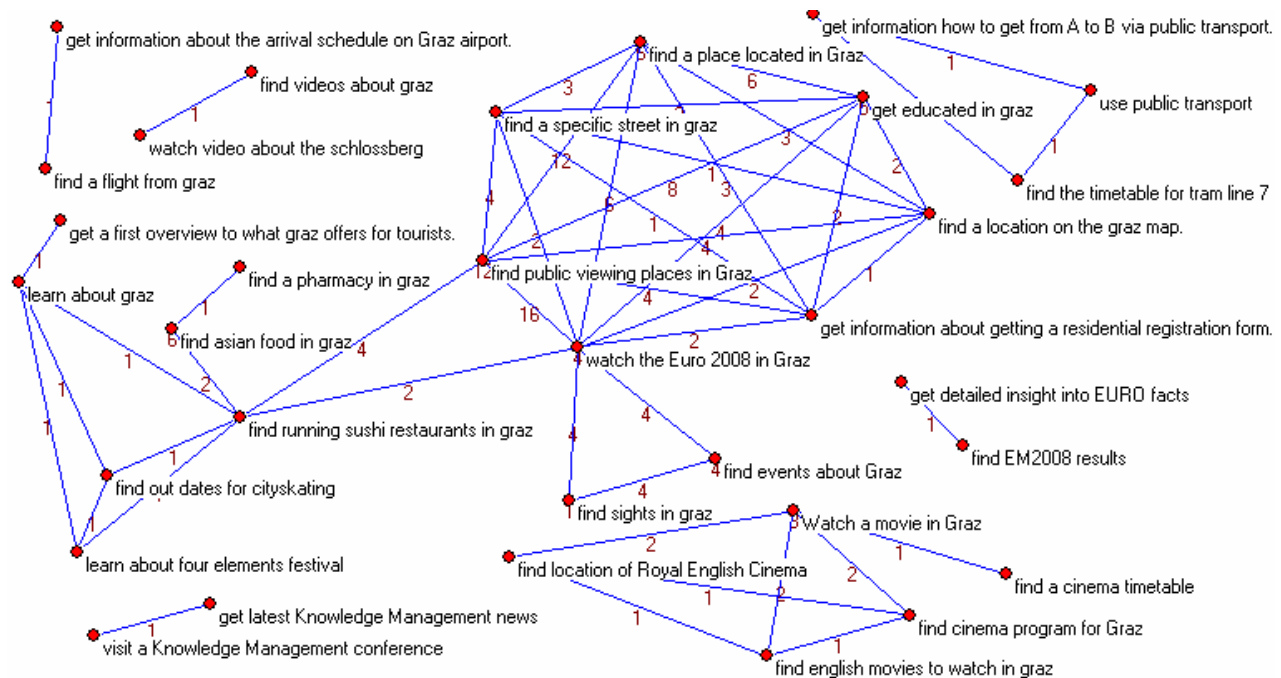


Figure 10 Purpose Tag Graph Obtained From Dataset A (Result of Folding the Bi-Partite OT Graph)

While this kind of analyses has been applied to study emergent semantics in the past [13], the graph illustrated in Figure 10 implies that purpose tags allow to study a different aspect of symbolic systems. Rather than studying emergent *semantics*, Figure 10 suggests that purpose tags allow to study emergent *pragmatics*, where the focus is on the usage context of resources rather than their meaning [17]. Using Pajek [14], the graph in Figure 10 was constructed by folding the bi-partite OT graph obtained from Dataset A after eliminating German tags, orphans and two specific websites that introduced noise (graz.at, google.com). In the folding process, only top-level domain information (e.g. tugraz.at) was retained.

Several points are interesting about this graph. First, recent research on search query log analysis has shown that search behaviour of users can be represented as a traversal of goal graphs [18]. Whereas traditional tag graphs represent the contents of the underlying resources, the purpose tag graph depicted in Figure 10 seems to be focused more on the goals of users. For example: in the tag graph above, the purpose tag “Watch the Euro 2008 in Graz” is strongly related to “find public viewing places in Graz” and “find a specific street in Graz”. Mapping search queries onto such goal graphs to refine or disambiguate search queries, and using these graphs to mine associations with resources could represent an interesting problem for future research. However, compared to traditional tag graphs, the graph in Figure 10 is rather sparse. This seems to be related to the specific interface chosen for the purpose tagging system, where users were encouraged to provide a single purpose tag per resource only. Overall though, the tag graph constructed in this study seems to generate interesting and meaningful relations between tags.

6. RELATION TO RELATED WORK

Purpose tags and purpose tag graphs could play a role in a number of current research areas and recent projects: *GOOSE*,

for example, is a prototypical goal-oriented search engine that aims to assist users in expressing their goals [10]. Purpose tags could be used to inform or refine queries in this context. *Miro*, another example, is an application that facilitates goal-oriented web browsing [4] by leveraging knowledge bases such as ConceptNet [11]. Knowledge about user goals represents a prerequisite for such applications to be effective. Another example, the work by Tanasescu et al [20], facilitates goal-oriented navigation of information through goal-driven service invocation - enabled by the Web Services Modeling Ontology WSMO [16], which is a web service description format that introduces the concept of user goals as a first class construct. Purpose tags and purpose tag graphs might have the potential to provide the background knowledge needed for such goal-oriented systems and applications.

7. CONCLUSIONS

This paper introduced purpose tags as a new kind of tags that has the potential to capture pragmatics (the usage context) of resources in social software applications rather than capturing their semantics. The case study we have conducted aimed to provide answers to 5 guiding research questions: 1. *Would users assign meaningful purpose tags?* While the concept of purpose tagging is new, we were able to show that users are comfortable assigning purpose tags in a case study involving a purpose tagging prototype. 2) *Do purpose tags accurately reflect plausible purposes of resources?* Based on a small sample, we could show that different subjects largely agree on whether purpose tags reflect the different purposes a given resource can serve accurately. 3) *Can purpose tagging improve search in social software?* While we were not able to quantitatively study this question in this work, goal-oriented social search seems to benefit from purpose tags as a mediator between user intent and the resources available in social software. 4) *Can purpose tags expand the vocabulary of existing tags?* Based on our

investigations, we could show that the vocabulary of purpose tags differs significantly from the vocabulary used in traditional kinds of tags. 5) *Are purpose tag graphs meaningful?* The tag graphs constructed from our purpose tagging prototype reveal pragmatic relations between different user goals and provide some interesting insights into motivational structures of a given user population. Future work focuses on further studies on purpose tagging, the identification of purpose tags from search query logs [19] and the construction of social-focused search algorithms that help guide user intent in social software applications.

8. ACKNOWLEDGMENTS

Thanks to A. Haselsberger und C. Ruggenthaler for implementing the purpose tagging prototype and to C. Körner for crawling the del.icio.us dataset. This work is funded by the FWF Austrian Science Fund Grant P20269 *TransAgere*. The Know-Center is funded within the Austrian COMET Program under the auspices of the Austrian Ministry of Transport, Innovation and Technology, the Austrian Ministry of Economics and Labor and by the State of Styria. COMET is managed by the Austrian Research Promotion Agency FFG.

9. REFERENCES

- [1] Ames, M. and Naaman, M. 2007. Why we tag: motivations for annotation in mobile and online media, in CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM, New York, NY, USA, 971-980.
- [2] Baeza-Yates, R., Calderón-Benavides, L. and González-Caro, C. 2006. The intention behind web queries, in F. Crestani, P. Ferragina and M. Sanderson, ed., Proceedings of String Processing and Information Retrieval (SPIRE), Springer, 98-109.
- [3] Broder, A. 2002. A taxonomy of web search, SIGIR Forum 36(2), 3-10.
- [4] Faaborg, A. and Lieberman, H. 2006. A goal-oriented web browser, in CHI '06: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM Press, New York, NY, USA, 751-760.
- [5] Golder, S. and Huberman, B. 2006. Usage patterns of collaborative tagging systems, Journal of Information Science 32(2), 198-208.
- [6] Heymann, P., Koutrika, G. and Garcia-Molina, H. 2008. Can social bookmarking improve web search?, in Proceedings of the International Conference on Web Search and Web Data Mining (WSDM'08), ACM, New York, NY, USA, 195-206.
- [7] Hotho, A., Jäschke, R., Schmitz, C. and Stumme, G. 2006. BibSonomy: A social bookmark and publication sharing system, in Proceedings of the Conceptual Structures Tool Interoperability Workshop at the 14th International Conference on Conceptual Structures, 87-102.
- [8] Krause, B., Hotho, A. and Stumme, G. 2008. A comparison of social bookmarking with traditional search, in Advances in Information Retrieval, 30th European Conference on IR Research, ECIR 2008, Springer, 101-113.
- [9] Lambiotte, R. and Ausloos, M. 2005. Collaborative tagging as a tripartite network, Arxiv preprint cs.DS/0512090 v2.
- [10] Liu, H., Lieberman, H. and Selker, T. 2002. GOOSE: A goal-oriented search engine with commonsense, in AH '02: Proceedings of the Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, Springer-Verlag, London, UK, 253-263.
- [11] Liu, H. and Singh, P. 2004. ConceptNet - Practical commonsense reasoning tool-kit, BT Technology Journal 22(4), 211-226.
- [12] Marlow, C., Naaman, M., Boyd, D. and Davis, M. 2006. HT06, tagging paper, taxonomy, Flickr, academic article, to read, in Hypertext '06: Proceedings of the Seventeenth Conference on Hypertext and Hypermedia, ACM, New York, NY, USA, 31-40.
- [13] Mika, P. 2005. Ontologies Are Us: A unified model of social networks and semantics, in Proceedings of the International Semantic Web Conference (ISWC'05), Springer, 522-536.
- [14] de Nooy, W., Mrvar, A. and Batagelj, V. 2005. Exploratory Social Network Analysis with Pajek, Cambridge University Press.
- [15] Norman, D. 1990. The design of everyday things, New York: Doubleday.
- [16] Roman, D., Keller, U., Lausen, H., de Bruijn, J., Lara, R., Stollberg, M., Polleres, A., Feier, C., Bussler, C. and Fensel, D. 2005. Web Service Modeling Ontology, Applied Ontology 1(1), 77-106.
- [17] Singh, M. 2002. The pragmatic web, IEEE Internet Computing 6(3), 4-5.
- [18] Strohmaier, M., Lux, M., Granitzer, M., Scheir, P., Liaskos, S. and Yu, E. 2007. How do users express goals on the web? - An exploration of intentional structures in web search, in Proceedings of the We-Know'07 International Workshop on Collaborative Knowledge Management for Web Information Systems, in conjunction with WISE'07, Nancy, France.
- [19] Strohmaier, M., Prettenhofer, P. and Lux M. 2008. Different Degrees of Explicitness in Intentional Artifacts - Studying User Goals in a Large Search Query Log. In Proceedings of the CSKGOI'08 Workshop on Commonsense Knowledge and Goal Oriented Interfaces, held in conjunction with IUT'08, Canary Islands, Spain.
- [20] Tanasescu, V., Gugliotta, A., Domingue, J., Villarias, L., Davies, R., Rowlett, M., Richardson, M. and Stincic, S. 2007. Geospatial data integration with semantic web services: the eMerges approach. In: The Geospatial Web, Springer Publishing.
- [21] Voss, J. 2007. Tagging, Folksonomy & co - renaissance of manual indexing?, arXiv:cs/0701072v2.
- [22] Yanbe, Y., Jatowt, A., Nakamura, S. and Tanaka, K. 2007. Can social bookmarking enhance search in the web?, in JCDL '07: Proceedings of the 7th ACM/IEEE-CS Joint Conference on Digital Libraries, ACM, New York, NY, USA, 107-116.